Event Rate Control in Online Advertising

To Appear in the Proceedings of the 2017 American Control Conference, Seattle, WA, May 24–26, 2017

Niklas Karlsson[1] and Qian Sang[2]

***Abstract*— Internet advertising is a relatively new area where feedback control has become critically important for scalable optimization. But using feedback control presents new challenges, one being the discontinuous nature of the input-output relationship of the plant to control. In this paper we propose an actuator and control algorithm for the specific objective of scalable event rate control in online advertising. The actuator makes the input-output relationship of the plant effectively continuous and with adjustable plant gain, and the feedback controller implements a PI controller to regulate the campaign-level event rate to stay at or above a reference value.**

I. INTRODUCTION

Advertising, which is a US$600 billion industry [1], has in recent years come to rely heavily on feedback control for online applications. 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 advertisements to Internet users). In short, each advertiser wishes to serve ads to those Internet users that generate the highest return on investment.

The allocation of ad impressions is handled in impression exchanges. Any advertiser may submit bids for any opportunity to show an ad, but only the highest bidder is awarded the impression. 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 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 more than ten years. See e.g. [2] for an early high-level introduction to the control problem and [3] 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 [4].

However, the problem considered in [4] is to maximize a value function given an ad budget. A slightly different problem is to control an average event rate, e.g., a campaign-level click-through or conversion rate, which is the focus of this paper. We are not aware of any previous attempt at solving this problem in the context of online advertising and restricted to decentralized (scalable) feedback control.

The paper is organized as follows. We define the control problem in Section II. By default the plant is discontinuous, but an actuation mechanism is proposed in Section III to effectively turn the input-output relationship of the plant continuous. In Section IV we describe how to model and tune the plant. The information is used to establish a nominal plant model that is used in Section V to design a feedback controller. In Section VI the control system is evaluated both in a simulated but realistic environment and on a real advertising campaign to assess performance and stability of the closed-loop control system. Finally, in Section VII we wrap up the paper with some concluding remarks and ideas of future work.

II. PROBLEM FORMULATION

The impression allocation for segment i is governed by a sealed second price auction [5], where b_i is the bid price submitted to the auction and a_i is the bid allocation, or the sampled fraction of auctions we choose to participate in. For problems related to maximization of value or return on investment, b_i is computed based on the estimated monetary value of impressions from segment i taking into account possible constraints on budget and return on investment [4]. However, sometimes an advertiser also cares about the average event rate, where an *event* is defined by the advertiser (e.g., a click or a conversion) and where the event rate p_i is defined as the probability that an impression from segment i turns into an event.

Suppose the campaign is submitting competitive bid prices (is the highest bidder) in segments labeled $i = 1, \dots, m$, and suppose the total number of available impressions in segment i is $n^{rel}_{avail,i}$. This paper deals with feedback-based control of a campaign's average event rate by way of adjusting a_i, and we neglect possible dynamic coupling between the computation of b_i and a_i.

The objective is to devise a feedback controller that adjusts a_i, $i = 1, \dots, m$, such that the average observed event rate of the campaign is at or above a prescribed reference value p^{ref}. We have access to site-level event rate estimates $\hat{p}_i \approx p_i$ and demand a computationally efficient (scalable) solution.

Scalability is obtained by the decoupled solution shown in the block diagram in Figure 1. Actuator is a static (memoryless) component processing the segment-level event rate estimates 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

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

[2]Q. Sang is Senior Research Scientist of R&D at Aol Platforms, 395 Page Mill Road, Palo Alto, CA 94306, USA qian.sang at teamaol.com

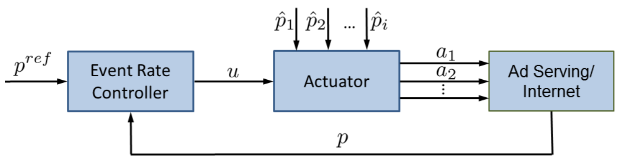

Fig. 1. Block diagram of the event rate control problem.

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.

III. 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 Figure 1). At our disposal are event rate estimates $\hat{p}_i$.

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 preferred the actuator is static (memory-less) and computationally inexpensive to use.

We impose the following requirements on the actuator mapping $a_i = g(\hat{p}_i, u)$, defined for $0 \leq \hat{p}_i \leq 1$ and $0 \leq u \leq 1$:

- g is static
- $0 \leq g(\hat{p}_i, u) \leq 1$ for all $\hat{p}_i, u$
- $g(\hat{p}_i, 0) = 1$ for all $\hat{p}_i$
- $g(\hat{p}_i, 1) = 0$ for $0 \leq \hat{p}_i < 1$
- $g(\hat{p}_i, u)$ is continous in $\hat{p}_i$ and u
- $g(\hat{p}_i, u)$ is decreasing in u for $0 < \hat{p}_i < 1$
- $g(\hat{p}_i, u)$ is increasing in $\hat{p}_i$ for $0 < u < 1$
- $g(\hat{p}_i, u)$ is a computationaly inexpensive mapping

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

A. Beta Distribution

The proposed actuator design makes use of properties of the so-called beta distribution from mathematical statistics [6]. The beta distribution with parameters α and β is a continuous probability distribution. If a random variable X follows the beta distribution, we write $X \sim \text{Beta}(\alpha, \beta)$. 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 variance σ^2 of X are

$$\mu := \text{E}(X) = \frac{\alpha}{\alpha+\beta},$$

$$\sigma^2 := \text{Var}(X) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}.$$

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

$$\alpha = \frac{\mu^2(1-\mu)}{\sigma^2} - \mu$$

$$\beta = (1-\mu)\left(\frac{\mu(1-\mu)}{\sigma^2} - 1\right)$$

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

B. 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 $0 \leq u \leq 1$. 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

$$\alpha_c(u) = cu$$

$$\beta_c(u) = c(1-u)$$

$$a_i = F(\hat{p}_i|\alpha_c(u), \beta_c(u))$$

if $0 < u < 1$; otherwise, $a_i = u$. The plots in Figure 2 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, and 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, we use $B(\hat{p}_i|\alpha,\beta)$ to denote the regularized

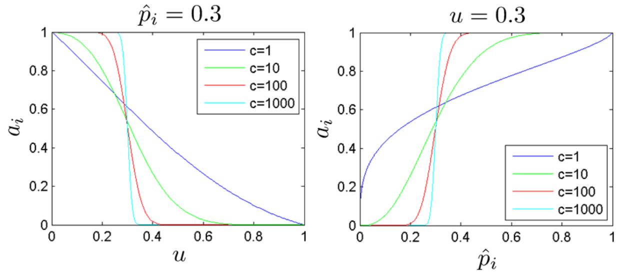

Fig. 2. The plots demonstrate how the bid allocation a_i for different values of c varies as a function of u for a fixed event rate $\hat{p}_i$ (left), and as a function of $\hat{p}_i$ for a fixed event rate u (right).

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 follows:

Algorithm 1 Beta actuation

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

Computation:

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

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.

Example 3.1: Figure 3 illustrates how the actuator re-

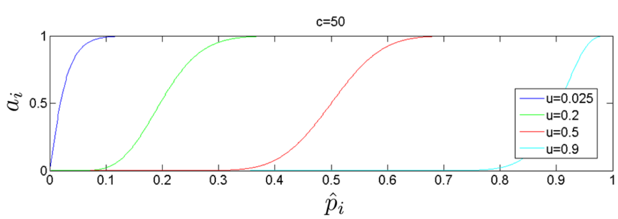

Fig. 3. The plot shows bid allocation a_i as a function of estimated event rate $\hat{p}_i$ for four different values of control signal u.

sponds 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 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 3.2: Figure 4 demonstrates how a_i varies as a

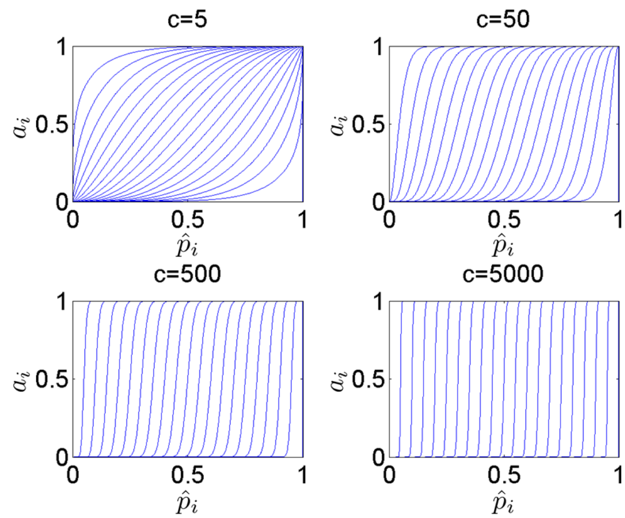

Fig. 4. The plots show bid allocation a_i as a function of estimated event rate $\hat{p}_i$ for $c = 5, 50, 500, 5000$, and for $u = 0, 0.05, 0.1, \ldots, 1$ (left to right curve in each plot).

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 $\mathbb{I}_{\{\hat{p}_i \geq u\}}$ when $c \to \infty$. ■

Example 3.3: Figure 5 shows an example of campaign-

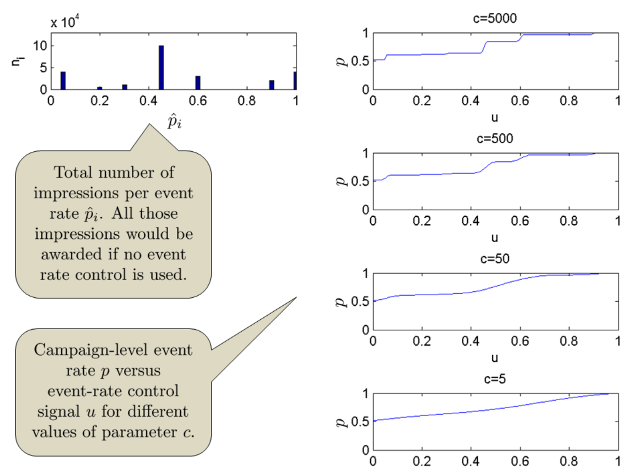

Fig. 5. Example of campaign-level relationship between control signal u and event rate p.

level relationship between control signal u and event rate p, depicted in the block diagram in Figure 1. 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 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 we regulate a_i 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. ■

IV. PLANT MODELING AND TUNING

In this section, we discuss plant modeling and tuning. The plant is defined by the mapping from campaign-level control input u to the campaign-level output p as shown in Fig. 1. The input-output relationship $u \to p$ may be tuned using the beta actuation parameter c.

For simplicity of presentation and without loss of generality, in the sequel of this paper we consider *in-view rate control* for display advertising. An impression is considered viewable if 50% of the ad pixels are in view for more than one second [7]. In the context of in-view rate control, an event is specifically an impression being 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 [7].

We first estimate the plant gain based on data from a population of 200 eCPM[1] advertising campaigns. Figure 6 shows the campaign-level in-view rate p vs. 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. 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 3.3. Note that smaller c leads to smoother slope curves, and its choice 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, we further generate the percentile plots in Fig. 7. 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. We choose $c = 50$, since it leads to a uniform plant

[1]An eCPM campaign is a campaign with an optimization objective of maximizing the total number of impressions for a given ad budget (eCPM=effective cost per thousand impressions).

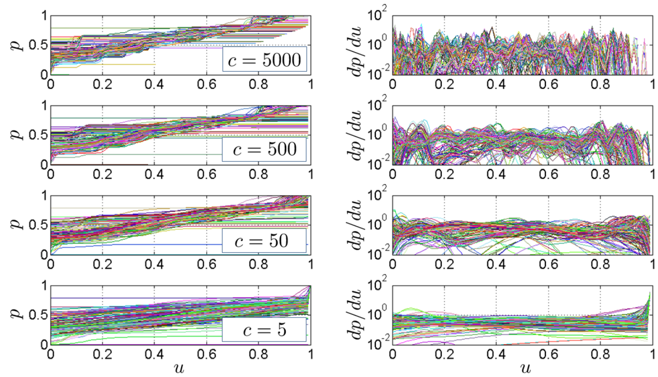

Fig. 6. In-view rate p (left) and in-view rate slope dp/du (right) vs. control signal u for select beta actuation configuration parameter c and for 200 representative ad campaigns.

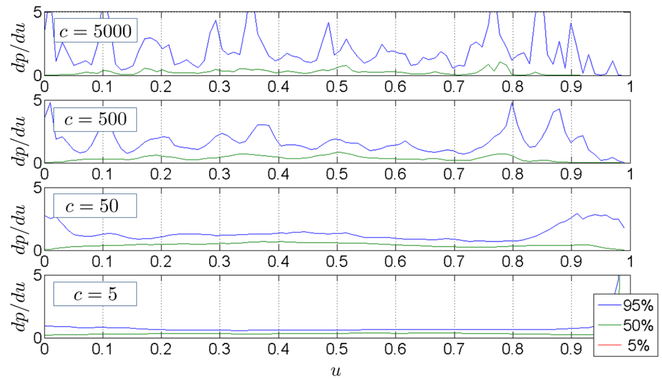

Fig. 7. Percentile plots of the in-view rate slope dp/du vs. control signal u for select Beta actuation sensitivity parameter c.

gain over a large range of the control signal u, e.g., for u in between roughly 0.05 and 0.83.

V. CONTROL DESIGN

We first present an in-view rate estimator that computes an estimate $\hat{p}$ of the campaign-level in-view rate p, as the feedback signal. A PI control scheme with windup protection is then employed for in-view rate control.

A. In-View Rate Estimator

Let $\{t_k\}, k = 0, 1, \ldots$, denote the sampling time instants and h the sampling period; and let $n_{meas}(t_k)$ and $n_{view}(t_k)$ denote the total (across all segments) number of measured impressions and the total number of viewable impressions, respectively, at time t_k. Let $\hat{p}(t_k)$ denote the campaign-level in-view rate estimate at time t_k. We compute $\hat{p}(t_k)$ from the impression counts as follows [8], [9]:

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

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 $n_{view}(t_k) \sim \text{Poisson}(n_{meas}(t_k)p)$ 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 [8], [9], [10].

B. In-View Rate Controller

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

We employ a PI controller with windup protection [11] to generate a control signal $u(t_k)$, 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 [11]:

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

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, we define $u_{low}(t_k)$ and $u_{high}(t_k)$ 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.

C. Selection of Design Parameters

The choice of design parameters is of significant importance to the overall control system performance. As can be seen from Fig. 7, the 95% curve with $c = 50$ provides a high 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, we opt for a 6dB ($\approx 20\log_{10} 2$) gain margin, 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 $T_{windup} < T_{int}$ [11]. We choose $T_{int}^{norm} = 3.33$ and $T_{windup}^{norm} = T_{int}^{norm}/1.05$, such that $h/T_{int} = 0.3$ and $T_{windup} = T_{int}/1.05$. Furthermore, we choose $\lambda = 0.9$ for the in-view rate estimator.

Table I summarizes the design parameter choices for the PI controller and in-view rate estimator.

TABLE I

SUMMARY OF DESIGN PARAMETERS.

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

VI. EXPERIMENT RESULTS

In this section, the performance of the proposed event rate control scheme has been evaluated in both a simulated environment and on a real advertising campaign.

A. Simulation Result

We first evaluate the proposed control system in a simulated environment. The plant is defined as a campaign with a total of $n_{avail}^{daily} = 2.4 \times 10^6$ 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 we draw a random number from $\text{Gamma}(\alpha, \beta)$ 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 $\text{Uniform}(0, 1)$ distribution.

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

TABLE II

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

The control performance is illustrated in Fig. 8 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, we are simulating a case in which the advertiser changes the in-view rate reference signal p^{ref}, as shown with a red dashed line in Fig. 8 (top). By computing u to drive the beta actuator, the proposed event rate controller regulates $\hat{p}$ to p^{ref}.

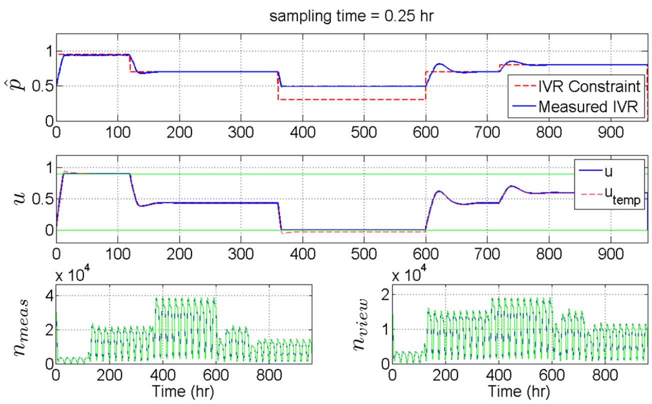

Fig. 8. Simulation results: campaign-level control system performance.

Note when p^{ref} is set high, e.g., $p^{ref} = 0.95$, during the first 120 hours, 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.

B. Performance on a Real Advertising Campaign

The proposed event rate control algorithm has been deployed to the AdLearn™ advertising optimization system developed by AOL. Figure 9 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 pacing and value maximization objective was fulfilled by a separate control algorithm. 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.

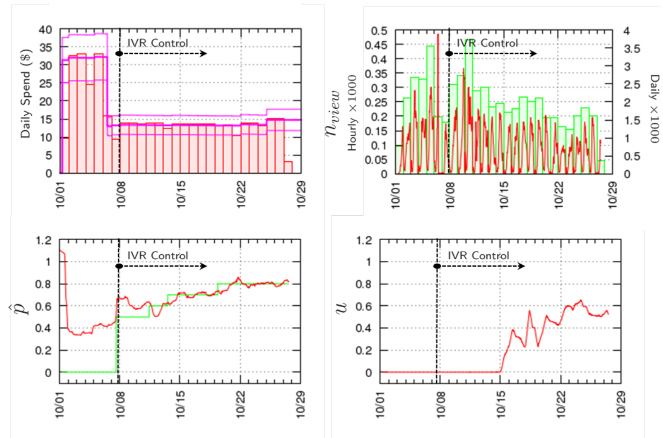

Fig. 9. Experiment results: campaign-level control system performance.

VII. CONCLUDING REMARKS

We have proposed an approach to actuation and feedback control of the average event rate of online advertising campaigns. In order to obtain a scalable solution the proposed system consists of a static actuator module consuming segment-level information, and a feedback controller consuming only campaign-level information. The challenge with this framework is that it may result in a plant with a discontinuous input-output relationship. The devised solution is a combination of an actuator that effectively turns the plant continuous and a PI controller that achieves reference tracking. The resulting control system has been evaluated on real advertising campaigns with excellent performance.

REFERENCES

[1] eMarketer. Advertisers will spend nearly $600 billion worldwide in 2015. *http://www.emarketer.com/Article/Advertisers-Will-Spend-Nearly-600-Billion-Worldwide-2015/1011691*, Last accesses Jan 6, 2016.
[2] Niklas Karlsson and Jianlong Zhang. Applications of feedback control in online advertising. *Proceedings of the 2013 American Control Conference*, pages 6008–6013, 2013.
[3] Niklas Karlsson. Adaptive control using Heisenberg bidding. *Proceedings of the 2014 American Control Conference*, pages 1304–1309, June 2014.
[4] Niklas Karlsson. Control problems in online advertising and benefits of randomized bidding strategies. *European Journal of Control*, 30:31–49, July 2016.
[5] Vijay Krishna. *Auction Theory*. Academic Press, 2002.
[6] George Casella and Roger L. Berger. *Statistical Inference*. Duxbury, second edition, 2001.
[7] IAB. State of viewability transaction 2015.
[8] Niklas Karlsson. Tutorial in Bayesian decision theory with applications in event rate estimation. Technical report, AOL Platforms R&D, September 2015.
[9] Niklas Karlsson. Adaptive estimation of small event rates. *Proceedings of the 55th IEEE Conference on Decision and Control*, pages 3732–3737, December 2016.
[10] James O. Berger. *Statistical Decision Theory and Bayesian Analysis, 2nd Edition*. Springer-Verlag, 1985.
[11] Karl Johan Åström and Tore Hägglund. *Advanced PID Control*. ISA – The Instrumentation, Systems, and Automation Society, 2005.