Scalable Multi-objective Optimization in Programmatic Advertising via Feedback Control

In the proceedings of the 60th IEEE Conf. on Decision and Control, Austin, Texas, USA, December 13-15, 2021

Niklas Karlsson[1]

Abstract— **The majority of online advertising is served through real-time bidding, and advertising campaigns are often defined as optimization problems. This paper deals with advertiser profit maximization subject to multiple advertiser performance constraints. The optimal bidding mechanism for a large family of multi-constrained advertising problems is derived, and it is demonstrated how the solution can be implemented as three separate subsystems; dealing with impression valuation, campaign control, and bid shading optimization, respectively. Feedback control plays a critical role to make this optimization scalable and adaptive. A proof of concept campaign control system is proposed and evaluated in simulations.**

I. Introduction

Optimal bidding for ad impressions is at the heart of online advertising, which is a large and rapidly growing industry. Feedback control is an essential part of the solution because of the extremely high dimensionality of the problem and the uncertain and dynamic nature of the environment.

A *Demand Side Platform* (DSP) serves as the middleman between an advertiser and one or more open exchange trading impressions, where an *impression* is the view of one ad. The DSP value proposition is to provide optimal bidding on behalf of the advertiser. It used to be enough to maximize the number of clicks or conversions while satisfying a spend constraint. However, the demands from advertisers keep increasing. Today advertisers often have additional constraints on their ad campaigns leading to a variety of multi-objective optimization problems, which is the subject of this paper.

Constrained optimization of display advertising in itself is not new. For example, optimal bidding for publisher revenue maximization in [1], [2] is defining a problem similar to what is considered in this paper. Their goal is to maximize the revenue subject to campaign budget and supply inventory constraints, but in their work the demand landscape is assumed partially or fully known. Strict advertiser optimization with limited knowledge of the competitive landscape is dealt with in [3], [4]; however, the solved problems involve particularly simple constraints and make significant assumptions on the environment.

A subproblem of ad optimization is to estimate event rates such as *click-through rate* (CTR) and *conversion rate* (CVR). This has received attention since the dawn of online advertising, and remains an active research area, e.g. [5], [6], [7] and many others. Another important and established subproblem is to adjust the bids based on error feedback to satisfy budget and performance constraints [8], [9], [10], [11], [12]. A third subproblem, which in recent time has gained attention due to an industry trend towards a first price cost model, is bid shading optimization [13], [14], [15], [16].

Our contribution is the optimal bidding mechanism for a large family of multi-constrained advertising problems under a first and second price cost model for the traded goods, and a proof of concept feedback control system to solve the campaign control subproblem of the overall optimization problem. It is demonstrated how the bid calculation lends itself to a decoupled implementation involving three separate and scalable systems; dealing with impression valuation, campaign control, and bid shading optimization, respectively. The bidding mechanism is derived using Lagrangian techniques together with probabilistic tools. The control system plays a critical role in the implementation of the optimization system, and enables a scalable and optimal solution. To prove the concept, the multiple-input, multiple-output control system is implemented as separate single-input, single-output pure integral (I) error feedback controllers and evaluated in a noise-free simulated scenario. The closed-loop performance is according to intuition and theory and serves as a benchmark test case.

The paper is organized as follows. Section II introduces key notation while Section III defines the problem that is solved later. Some definitions and preliminary results are provided in Section IV. The key contribution and the optimal bidding mechanism is presented in Section V. Section VI elaborates on how to implement the solution with a feedback controller. The result is made clear with help of an example in Section VII, and conclusions plus ideas of future work are discussed in Section VIII.

II. Set-up and Notation

The goal is to solve a constrained optimization problem via real time bidding on impressions on behalf of one specific ad campaign. Each impression is awarded to the highest bidder, and the winning bidder is charged according to a first or second price cost model [17]. The cost model is known before the bids are submitted, and for a first price impression the winner pays an amount equal to its own bid, whereas for a second price impression the winner pays an amount equal to the second highest bid. Notation and assumptions used extensively in the paper are listed below:

- The set of all impression opportunities is denoted Ω, and subsets Ω_1 and Ω_2 are impression opportunities sold based on a first and second price cost model, respectively. Let $\Omega_1 \cap \Omega_2 = \emptyset$ and $\Omega_1 \cup \Omega_2 = \Omega$.

[1]N. Karlsson is Chief Scientist and Vice President of Demand Platforms R&D, Yahoo Research, 1199 Coleman Ave San Jose, CA 95110, USA nkarlsson8 at gmail.com

- The uppercase of a Roman letter denotes a random variable, while a lowercase represents its expected value (future event) or realization (historical event).
- The *bid price* b_i, $i \in \Omega$, is the bid amount submitted to the auction for the ith impression opportunity on behalf of the represented advertiser, and is a decision variable of the optimization problem.
- The *context* of impression opportunity i is information available to the bidder to be used to calculate a bid. It may include e.g. the website and placement of the ad, the type of devise where the impression request origins, and various demographic information about the user.
- The *highest competing bid* of impression i is a random variable $B_i^\star \in \mathbb{R}_{>0}$, and $b_i^\star$ denotes its realized value.
- The *cumulative distribution function* (CDF) of $B_i^\star$, given the context, is $F_{B_i^\star}(b)$, while the *probability density function* (PDF) is $f_{B_i^\star}(b) = dF_{B_i^\star}(b)/db$, which is assumed to be continuous. See Figure 1 for a few examples of CDFs.

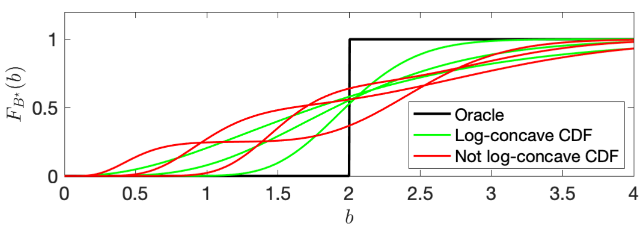

Fig. 1. Seven sample CDFs: 1) The black curve is available to an oracle and indicates a perfect a priori knowledge of the highest competing bid, 2) the green curves are log-concave functions with nice properties that are exploited later, and 3) the red curves are general not log-concave functions.

- The ith impression is awarded to us if $b_i \geq b_i^\star$
- The *ad cost* C_i of an awarded impression i is

$$C_i = b_i \mathbb{I}_{\{i \in \Omega_1\}} + B_i^\star \mathbb{I}_{\{i \in \Omega_2\}}, \quad (1)$$

  where $\mathbb{I}_A$ is the indicator function satisfying $\mathbb{I}_A = 1$, if $A =$ true, and $\mathbb{I}_A = 0$, otherwise. Note, Ω_2 is the complement of Ω_1, hence $C_i = b_i$ or $C_i = B_i^\star$.
- The *event count* $N_{j,i} \in \mathbb{N}$ is the number of events of type j following an awarded impression. Examples of event types are impressions, clicks, and conversions (product sales). Furthermore, the *impression value* $V_i \in \mathbb{R}_{>0}$ is the advertising value attributed to the ith impression, if awarded. It encodes the branding and/or performance value, and is typically a function of one or more of the event counts.
- The *total number of j-events*, N_j, the *total value*, V, and the *cumulative ad cost*, C, are

$$N_j = \sum_{i \in \Omega} N_{j,i} \mathbb{I}_{\{b_i \geq B_i^\star\}}, \quad j = 1, \dots, m, \quad (2)$$

$$V = \sum_{i \in \Omega} V_i \mathbb{I}_{\{b_i \geq B_i^\star\}}, \quad \text{and} \quad (3)$$

$$C = \sum_{i \in \Omega_1} b_i \mathbb{I}_{\{b_i \geq B_i^\star\}} + \sum_{i \in \Omega_2} B_i^\star \mathbb{I}_{\{b_i \geq B_i^\star\}}. \quad (4)$$

- The total profit is the difference between total value and cumulative ad cost, hence, equals $V - C$.
- The expected impression value and event count, given the context, are $\mathrm{E}V_i = v_i$ and $\mathrm{E}N_{j,i} = p_{j,i}$ (breaking from the convention that suggests the notation $n_{j,i}$). Event type $j = 1$ is without loss of generality an impression, which implies that $p_{1,i} = 1$, for all $i \in \Omega$.
- The event count $N_{j,i}$ and $B_i^\star$, as well as, the impression value V_i and $B_i^\star$ are assumed conditionally independent, given the context of the impression.

Finally, without explicitly mentioning it later, all expected values throughout the paper are conditioned on the context.

III. Problem Formulation

The objective is to compute the bids b_i, for all $i \in \Omega$, that maximize the expected total profit, $\mathrm{E}V - \mathrm{E}C$, subject to constraints on total spend, *effective cost per event* (eCPX), and *events per impression rate* (ER). Constraints are defined by non-negative parameters ξ_1, $\xi_{2,j}$, $\xi_{3,j}$ and the problem is mathematically defined by

$$\underset{\{b_i \in \mathbb{R} | \forall i \in \Omega\}}{\text{maximize}} \quad \mathrm{E}V - \mathrm{E}C \quad (5)$$

subject to

$$\begin{aligned} \mathrm{E}C &\leq \xi_1 && \text{(spend)} && (6) \\ \mathrm{E}C &\leq \xi_{2,j} \mathrm{E}N_j && (\text{eCPX}, j = 1, \dots, m) && (7) \\ \xi_{3,j} \mathrm{E}N_1 &\leq \mathrm{E}N_j && (\text{ER}, j = 2, \dots, m) && (8) \end{aligned}$$

The cardinality of Ω is in the order of millions or billions, which makes the problem extremely high-dimensional. Moreover, the cardinality of Ω, the impression values v_i, the event rates $p_{j,i}$, and the competing bid prices $b_i^\star$ are a priori unknown. Finally, the impression opportunities occur over time, not all at once, which is exploited in Section VI. Typically, an advertiser prescribes the constraints on a per day basis (a spend budget ξ_1 per day, etc.)

IV. Preliminaries

Define *constraint vector* $\bar{\xi}$, *impression value vector* $\bar{v}_i$, *Lagrange multiplier vector* $\bar{\lambda}$, and *control signal vector* $\bar{u}$ as

$$\begin{aligned} \bar{\xi} &:= \begin{bmatrix} \xi_1 & \xi_{2,1} & \cdots & \xi_{2,m} & \xi_{3,2} & \cdots & \xi_{3,m} \end{bmatrix}^T, \\ \bar{v}_i &:= \begin{bmatrix} v_i & p_{2,i} & p_{3,i} & \cdots & p_{m,i} \end{bmatrix}^T, \\ \bar{\lambda} &:= \begin{bmatrix} \lambda_1 & \lambda_{2,1} & \cdots & \lambda_{2,m} & \lambda_{3,2} & \cdots & \lambda_{3,m} \end{bmatrix}^T, \\ \bar{u} &:= \begin{bmatrix} u_0 & u_v & u_2 & \cdots & u_j & \cdots & u_m \end{bmatrix}^T, \end{aligned}$$

where the elements of $\bar{\lambda}$ and $\bar{u}$ are defined in Section V.

Vector $\bar{\xi}$ contains the constraints imposed by the advertiser. These are typically entered by humans and change infrequently. They are known and can often be viewed as constants. Vector $\bar{v}_i$ consists of quantities such as impression value, click-through-rate, and various conversion-rates. These are typically known only as estimates and differ widely from one impression opportunity i to the next. Furthermore, vector $\bar{\lambda}$ is composed of multipliers, which in an implementation are used as tuning knobs and adjusted gracefully over time toward their optimal values for which the constraints are satisfied. Finally, vector $\bar{u}$ is a rational static function of $\bar{\xi}$ and $\bar{\lambda}$ that facilitates an efficient real time bid computation.

The following results are used later.

Lemma 4.1:

$$\mathrm{E}(\mathbb{I}_{\{b_i \geq B_i^\star\}}) = F_{B_i^\star}(b_i) \quad (9)$$

$$\mathrm{E}\big(B_i^\star \mathbb{I}_{\{b_i \geq B_i^\star\}}\big) = b_i F_{B_i^\star}(b_i) - \int_0^{b_i} F_{B_i^\star}(z)dz \quad (10)$$

Proof: (I) Identity (9) is obtained from $\mathrm{E}(\mathbb{I}_{\{b_i \geq B_i^\star\}}) = \int_{-\infty}^{\infty} \mathbb{I}_{\{b_i \geq z\}} f_{B_i^\star}(z)dz = \int_{-\infty}^{b_i} f_{B_i^\star}(z)dz = F_{B_i^\star}(b_i)$.

(II) Identity (10) is obtained via integration by parts from $\mathrm{E}\left(B_i^\star \mathbb{I}_{\{b_i \geq B_i^\star\}}\right) = \int_{-\infty}^{\infty} z \mathbb{I}_{\{b_i \geq z\}} f_{B_i^\star}(z)dz = \int_0^{b_i} z f_{B_i^\star}(z)dz = b_i F_{B_i^\star}(b_i) - \int_0^{b_i} F_{B_i^\star}(z)dz$. ■

V. Optimal Bidding Mechanism

The scale of the problem, the uncertain competitive landscape, and the a priori unknown impression value and event rates of each impression opportunity, makes it virtually impossible to solve the optimization problem using a plan-based algorithm. Instead a Lagrangian approach is adopted to rewrite the problem as three largely independent subproblems that are solved with scalable and adaptive methods.

The following theorem is the main result of the paper and states the necessary conditions for optimality.

Theorem 5.1: The optimal bid b_i^{opt}, for all $i \in \Omega$, satisfies

$$b_i^{opt} = \begin{cases} \underset{b\in[0,b_i^u]}{\operatorname{argmax}} (b_i^u - b)F_{B_i^\star}(b), & i \in \Omega_1, \\ b_i^u, & i \in \Omega_2, \end{cases}$$

whenever $b_i^u \geq 0$; and $b_i^{opt} = 0$, if $b_i^u < 0$. The *adjusted impression value* b_i^u is

$$b_i^u = \frac{v_i + \sum_{j=1}^m p_{j,i}\xi_{2,j}\lambda_{2,j} + \sum_{j=2}^m (p_{j,i} - \xi_{3,j})\lambda_{3,j}}{1 + \lambda_1 + \sum_{j=1}^m \lambda_{2,j}}, \quad (11)$$

and where inequalities (6)-(8) hold, as well as, $\lambda_1(\mathrm{E}C - \xi_1) = \lambda_{2,j}(\mathrm{E}C - \xi_{2,j}\mathrm{E}N_j) = \lambda_{3,j}(\xi_{3,j}\mathrm{E}N_1 - \mathrm{E}N_j) = 0$.

Proof: The Lagrangian of (5)-(8) is

$$\mathcal{L} = \mathrm{E}V - \mathrm{E}C - \lambda_1(\mathrm{E}C - \xi_1) - \sum_{j=1}^m \lambda_{2,j}(\mathrm{E}C - \xi_{2,j}\mathrm{E}N_j) - \sum_{j=2}^m \lambda_{3,j}(\xi_{3,j}\mathrm{E}N_1 - \mathrm{E}N_j). \quad (12)$$

If there exists b_i, $\forall i \in \Omega$, and $\lambda_1, \lambda_{2,j}, \lambda_{3,j} \geq 0$, $\forall j$, such that the b_i's maximize $\mathcal{L}$, the inequalities (6)-(8) are satisfied, and $\lambda_1(\mathrm{E}C - \xi_1) = \lambda_{2,j}(\mathrm{E}C - \xi_{2,j}\mathrm{E}N_j) = \lambda_{3,j}(\xi_{3,j}\mathrm{E}N_1 - \mathrm{E}N_j) = 0$; then, due to the Lagrangian sufficiency theorem [18], these values of b_i solve (5)-(8).

For fixed values of λ_1, $\lambda_{2,j}$, and $\lambda_{3,j}$, for all j, we compute the optimal values of b_i. Collect all terms of $\mathrm{E}V$, $\mathrm{E}C$, $\mathrm{E}N_1$, and $\mathrm{E}N_j$ in (12), and rewrite the Lagrangian as

$$\mathcal{L} = \mathrm{E}V - \left(1 + \lambda_1 + \sum_{j=1}^m \lambda_{2,j}\right)\mathrm{E}C + \lambda_1\xi_1 + \left(\lambda_{2,1}\xi_{2,1} - \sum_{j=2}^m \lambda_{3,j}\xi_{3,j}\right)\mathrm{E}N_1 + \sum_{j=2}^m (\lambda_{2,j}\xi_{2,j} + \lambda_{3,j})\mathrm{E}N_j. \quad (13)$$

Random variables $N_{j,i}$ and V_i are by assumption (see Section II) conditionally independent of $B_i^\star$, given the context; hence, $\mathrm{E}(N_{j,i}\mathbb{I}_{\{b_i \geq B_i^\star\}}) = E(N_{j,i})E(\mathbb{I}_{\{b_i \geq B_i^\star\}})$ and $\mathrm{E}(V_i\mathbb{I}_{\{b_i \geq B_i^\star\}}) = E(V_i)E(\mathbb{I}_{\{b_i \geq B_i^\star\}})$. Since $\mathrm{E}N_{j,i} = p_{j,i}$ and $\mathrm{E}V_i = v_i$, it follows from (2)-(4) that

$$\mathrm{E}N_j = \sum_{i\in\Omega} p_{j,i}\mathrm{E}\left(\mathbb{I}_{\{b_i \geq B_i^\star\}}\right),$$

$$\mathrm{E}V = \sum_{i\in\Omega} v_i\mathrm{E}\left(\mathbb{I}_{\{b_i \geq B_i^\star\}}\right),$$

$$\mathrm{E}C = \sum_{i\in\Omega_1} b_i\mathrm{E}\left(\mathbb{I}_{\{b_i \geq B_i^\star\}}\right) + \sum_{i\in\Omega_2} \mathrm{E}\left(B_i^\star\mathbb{I}_{\{b_i \geq B_i^\star\}}\right).$$

Use Lemma 4.1 to rewrite these expected values as

$$\mathrm{E}N_j = \sum_{i\in\Omega} p_{j,i}F_{B_i^\star}(b_i),$$

$$\mathrm{E}V = \sum_{i\in\Omega} v_iF_{B_i^\star}(b_i),$$

$$\mathrm{E}C = \sum_{i\in\Omega} b_iF_{B_i^\star}(b_i) - \sum_{i\in\Omega_2}\int_0^{b_i} F_{B_i^\star}(z)dz.$$

Plug in the derived expressions for $\mathrm{E}N_j$, $\mathrm{E}V$, and $\mathrm{E}C$ into the Lagrangian (13), and gather all terms of $F_{B_i^\star}(b_i)$ and $\int_0^{b_i} F_{B_i^\star}(z)dz$. We obtain

$$\mathcal{L} = \sum_{i\in\Omega}\left(v_i + \sum_{j=1}^m p_{j,i}\xi_{2,j}\lambda_{2,j} + \sum_{j=2}^m (p_{j,i} - \xi_{3,j})\lambda_{3,j} - \left(1 + \lambda_1 + \sum_{j=1}^m \lambda_{2,j}\right)b_i\right)F_{B_i^\star}(b_i) + \left(1 + \lambda_1 + \sum_{j=1}^m \lambda_{2,j}\right)\sum_{i\in\Omega_2}\int_0^{b_i} F_{B_i^\star}(z)dz + \lambda_1\xi_1. \quad (14)$$

To save space, introduce helper functions $g_0(\bar{v}_i, \bar{\xi}, \bar{\lambda})$ and $g_1(\bar{\lambda})$ defined by

$$g_0(\bar{v}_i, \bar{\xi}, \bar{\lambda}) = v_i + \sum_{j=1}^m p_{j,i}\xi_{2,j}\lambda_{2,j} + \sum_{j=2}^m (p_{j,i} - \xi_{3,j})\lambda_{3,j} \quad (15)$$

and

$$g_1(\bar{\lambda}) = 1 + \lambda_1 + \sum_{j=1}^m \lambda_{2,j}. \quad (16)$$

Substitute for $g_0(\bar{v}_i, \bar{\xi}, \bar{\lambda})$ and $g_1(\bar{\lambda})$ in (14) to obtain

$$\mathcal{L} = \sum_{i\in\Omega}\left(g_0(\bar{v}_i, \bar{\xi}, \bar{\lambda}) - g_1(\bar{\lambda})b_i\right)F_{B_i^\star}(b_i) + g_1(\bar{\lambda})\sum_{i\in\Omega_2}\int_0^{b_i} F_{B_i^\star}(z)dz + \lambda_1\xi_1. \quad (17)$$

All Lagrange multipliers are non-negative, hence $g_1(\bar{\lambda}) > 0$. Moreover, the adjusted impression value, as defined in (11), can be written as $b_i^u = g_0(\bar{v}_i, \bar{\xi}, \bar{\lambda})/g_1(\bar{\lambda})$. A simple rearrangement of (17) then yields

$$\mathcal{L} = g_1(\bar{\lambda})\sum_{i\in\Omega}\left(\left(b_i^u - b_i\right)F_{B_i^\star}(b_i) + \int_0^{b_i} F_{B_i^\star}(z)dz\,\mathbb{I}_{\{i\in\Omega_2\}}\right) + \lambda_1\xi_1.$$

Define the expression inside the big parenthesis as the expected *surplus per bid response*, $s_i(b_i^u, b_i)$; i.e.,

$$s_i(b_i^u, b_i) = (b_i^u - b_i)F_{B_i^\star}(b_i) + \int_0^{b_i} F_{B_i^\star}(z)dz\mathbb{I}_{\{i\in\Omega_2\}}. \quad (18)$$

Therefore

$$\mathcal{L} = g_1(\bar\lambda)\sum_{i\in\Omega} s_i(b_i^u, b_i) + \lambda_1\xi_1. \quad (19)$$

Since $g_1(\bar\lambda) > 0$, the bids that maximize $s_i(b_i^u, b_i)$ individually for each i also maximize $\mathcal{L}$. The partial derivative of $s_i(b_i^u, b_i)$ with respect to b_i is obtained from (18) as

$$\frac{\partial s_i(b_i^u, b_i)}{\partial b_i} = (b_i^u - b_i)f_{B_i^\star}(b_i) - F_{B_i^\star}(b_i)\mathbb{I}_{\{i\in\Omega_1\}}, \quad (20)$$

where the relationship $\mathbb{I}_{\{i\in\Omega_2\}} = 1 - \mathbb{I}_{\{i\in\Omega_1\}}$ has been used. Both $f_{B_i^\star}(b_i)$ and $F_{B_i^\star}(b_i)$ are zero if $b_i \le 0$, hence, if $b_i^u \le 0$, then $s_i(b_i^u, b_i)$ is maximized (zero) at $b_i^{opt} = 0$. Otherwise, if $b_i^u > 0$, consider cases $i \in \Omega_1$ and $i \in \Omega_2$, separately.

Case $\mathbf{i \in \Omega_1}$: It follows from (20) that

$$\frac{\partial s_i(b_i^u, b_i)}{\partial b_i} \begin{cases} = 0, & b_i < 0, \\ \text{depends on } F_{B_i^\star}(b), & 0 \le b_i \le b_i^u, \\ \le 0, & b_i > b_i^u. \end{cases} \quad (21)$$

Typically, no closed form expression for b_i^{opt} exists if $i \in \Omega_1$; however, it is clear from (18) that $s_i(b_i^u, b_i)$ is zero if $b_i \le 0$, non-negative if $b_i \in [0, b_i^u]$, and non-positive if $b_i > b_i^u$. Hence, $s_i(b_i^u, b_i)$ is maximized for a value of b_i in the interval $[0, b_i^u]$. The optimal bid satisfies

$$b_i^{opt} = \underset{b\in[0,b_i^u]}{\text{argmax}}(b_i^u - b)F_{B_i^\star}(b), \quad i \in \Omega_1. \quad (22)$$

Case $\mathbf{i \in \Omega_2}$: It follows from (20) that

$$\frac{\partial s_i(b_i^u, b_i)}{\partial b_i} = (b_i^u - b_i)f_{B_i^\star}(b_i) \begin{cases} \ge 0, & b_i \le b_i^u, \\ \le 0, & b_i > b_i^u. \end{cases} \quad (23)$$

Hence, $s_i(b_i^u, b_i)$ is nondecreasing in b_i if $b_i \le b_i^u$, and is nonincreasing otherwise. Therefore, $s_i(b_i^u, b_i)$ is maximized in regards to b_i by choosing

$$b_i^{opt} = b_i^u, \quad i \in \Omega_2, \quad (24)$$

which completes the proof. ■

Theorem 5.1 suggests that the optimization problem can be decomposed into three subproblems, representing a three player non-cooperative game:

1) **Impression valuation:** Estimate the impression value vector $\bar v_i$ for each impression opportunity.
2) **Campaign control:** Estimate the impression-independent optimal Lagrange multiplier vector $\bar\lambda$.
3) **Bid Shading Optimization:** Estimate the CDF $F_{B_i^\star}(b)$ and compute b_i^{opt} based on $b_i^u = b_i^u(\bar v_i, \bar\lambda)$ and $F_{B_i^\star}(b)$.

In practice, $\bar\lambda$ is updated at discrete time points, such as once every five minutes or one hour, while b_i^u is computed for each individual impression opportunity. The computations of b_i^u and b_i^{opt} are subject to tight time constraints, and must be completed within a few milliseconds. The following corollary provides the means of making the real-time computation of b_i^u efficient.

Corollary 5.1: The adjusted impression value b_i^u in Theorem 5.1 can be expressed

$$b_i^u = [1, \bar v_i^T]\bar u,$$

where the elements of $\bar u := [u_0, u_v, u_2, \ldots, u_j, \ldots, u_m]^T$ are

$$u_0 = \frac{\lambda_{2,1}\xi_{2,1} - \sum_{j=2}^m \lambda_{3,j}\xi_{3,j}}{1 + \lambda_1 + \sum_{\ell=1}^m \lambda_{2,\ell}},$$
$$u_v = \frac{1}{1 + \lambda_1 + \sum_{\ell=1}^m \lambda_{2,\ell}},$$
$$u_j = \frac{\lambda_{2,j}\xi_{2,j} + \lambda_{3,j}}{1 + \lambda_1 + \sum_{\ell=1}^m \lambda_{2,\ell}}, \quad j = 2, \ldots, m.$$

Proof: The result is obtained from (11) via a trivial rearrangement of the numerator. ■

Figure 2 provides a block diagram of the modularized

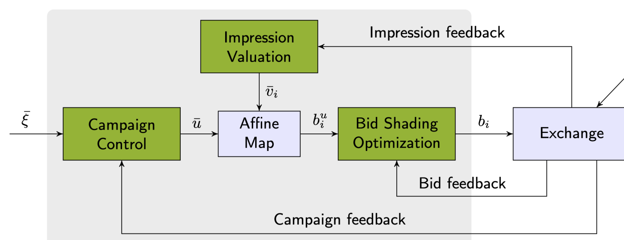

Fig. 2. The interconnected optimization system is enclosed in the large grey rectangle with dynamic subsystems marked green. Campaign-level signals are represented by $\bar\xi \in \mathbb{R}^{2m}$ and $\bar u \in \mathbb{R}^{m+1}$, while signals $\bar v_i \in \mathbb{R}^m$, and $b_i^u, b_i, b_i^\star \in \mathbb{R}$ are defined for each individual impression opportunity.

system. The benefit of the decomposition is best appreciated by recognizing that the impression opportunities in Ω do not arrive all at once, but as a time series throughout the campaign flight. Moreover, impression value vector $\bar v_i$, control signal vector $\bar u$ (via vectors $\bar\lambda$ and $\bar\xi$), and CDF $F_{B_i^\star}(b_i)$ are unknown a priori and must be estimated online.

The output $\bar u$ of *campaign control* is combined, for each impression opportunity i, with an output $\bar v_i$ of *impression valuation* via an affine dot product $b_i^u = [1, \bar v_i^T]\bar u$. If $b_i^u \ge 0$, then this adjusted impression value is used in *bid shading optimization* to compute the final bid $b_i = \arg\max_{b\in[0,b_i^u]} s_i(b_i^u, b)$, which is submitted to the auction.

The following two theorems capture important properties of the bid shading subproblem that relates b_i^{opt} and b_i^u.

Theorem 5.2: The expected surplus per bid request maximizing bid $b_i^{opt}(b_i^u) := \arg\max_{b\in[0,b_i^u]} s(b_i^u, b)$ is a non-decreasing function of the adjusted impression value b_i^u.

Proof: To avoid clutter, index i, identifying a specific impression opportunity, is omitted in the proof. Consider cases $i \in \Omega_1$ and $i \in \Omega_2$, separately.

Case $\mathbf{i \in \Omega_1}$: Consider two different adjusted impression values $b_a^u, b_b^u \in \mathbb{R}$. It follows from (18) that the surplus $s(b^u, b)$ is an affine function of b^u, and that

$$s(b_b^u, b^{opt}(b_a^u)) = s(b_a^u, b^{opt}(b_a^u)) + (b_b^u - b_a^u)F_{B^\star}(b^{opt}(b_a^u)),$$

$$s(b_b^u, b) = s(b_a^u, b) + (b_b^u - b_a^u)F_{B^\star}(b),$$

for all $b \in \mathbb{R}$. Assume without loss of generality that $b_b^u > b_a^u$ and let $b^{opt}(b_a^u) > b$. Then

$$\begin{aligned} s(b_b^u, b^{opt}(b_a^u)) - s(b_b^u, b) &= s(b_a^u, b^{opt}(b_a^u)) - s(b_a^u, b) \\ &+ (b_b^u - b_a^u)\big(F_{B^\star}(b^{opt}(b_a^u)) - F_{B^\star}(b)\big). \end{aligned} \quad (25)$$

Consider the right hand side of (25). By definition, the expected surplus per bid request maximizing bid for an impression opportunity with adjusted impression value b_a^u is $b^{opt}(b_a^u)$; i.e., no bid b generates a larger surplus when the adjusted impression value is b_a^u. Hence

$$s(b_a^u, b^{opt}(b_a^u)) - s(b_a^u, b) \quad \geq \quad 0.$$

Moreover, by assumption, $b_b^u - b_a^u \geq 0$, and by virtue of $F_{B^\star}(b)$ being a CDF, it is a non-negative and non-decreasing function of b, and since $b^{opt}(b_a^u) > b$, it follows that

$$F_{B^\star}(b^{opt}(b_a^u)) - F_{B^\star}(b) \quad \geq \quad 0.$$

Hence, the right hand side of (25) must be non-negative; i.e.,

$$s(b_b^u, b^{opt}(b_a^u)) - s(b_b^u, b) \quad \geq \quad 0, \qquad b^{opt}(b_a^u) > b.$$

In other words, when the adjusted impression value is b_b^u ($> b_a^u$), then the surplus per bid request at bid $b^{opt}(b_a^u)$ is no smaller than the surplus per bid request at any bid $b < b^{opt}(b_a^u)$. As a consequence, the surplus per bid request maximizing bid for the impression opportunity with adjusted impression value $b_b^u > b_a^u$ must satisfy

$$b^{opt}(b_b^u) \quad \geq \quad b^{opt}(b_a^u).$$

Case $\mathbf{i} \in \mathbf{\Omega_2}$: Due to Theorem 5.1, $b^{opt} = b^u$, hence, b^{opt} is non-decreasing in b^u, which completes the proof. ∎

Theorem 5.3: If $F_{B_i^\star}(b)$ is log-concave in the interval $(0, b_i^u)$, and $b_i^u > 0$, then any local maximum of $s(b_i^u, b)$ relative $b \in (0, b_i^u)$ is a global maximum and an optimal bid.

Proof: To reduce clutter, index i is omitted in the proof.

Case $\mathbf{i} \in \mathbf{\Omega_1}$: Recall from the proof of Theorem 5.1 that $s(b^u, b) = (b^u - b)F_{B^\star}(b)$ and that the maximizing b lies between 0 and b^u. Since $F_{B^\star}(b)$ is log-concave in this interval, $\log(s(b^u, b)) = \log(b^u - b) + \log(F_{B^\star}(b))$ is defined in $(0, b^u)$ and the sum of two concave functions. This implies that $\log(s(b^u, b))$ is concave in the interval, and any local maximum of $s(b^u, b)$ in the interval is a global maximum and an optimal bid.

Case $\mathbf{i} \in \mathbf{\Omega_2}$: The result is trivial since in this case $\arg\max_b s(b^u, b) = b^u$ and independent of $F_{B^\star}(b)$. ∎

Many common probability distributions have log-concave CDFs. This includes, but is not limited to distributions such as gamma, log-normal, normal, and logistic. If the context of served impressions is sufficiently granular, then one can typically approximate $F_{B_i^\star}(b)$ by a log-concave function (see Figure 1), which makes bid shading optimization easier.

VI. Application

The complete solution consists of impression valuation, campaign control, and bid shading optimization (see Figure 2). Various embodiments of impression valuation [5], [6], [7] and bid shading optimization [13], [14], [15], [16] are readily available in the literature, and the remainder of this article deals with a proof of concept implementation of campaign controller.

A. Control Problem and Plant Modeling

Consider a time-sampled implementation of the system defined by equidistant time points indexed $t = 1, 2, \ldots, t_{max}$; where t_{max} is the number of time intervals in the optimization window (e.g. one day). Let $\Omega(t)$, $v(t)$, $c(t)$, and $n_j(t)$ denote the portion of Ω, v, c, and n_j realized in time interval t. Hence, $\Omega = \sum_t \Omega(t)$, $v = \sum_t v(t)$, $c = \sum_t c(t)$, and $n_j = \sum_t n_j(t)$; where all summations are for t from 1 to t_{max}. The values of $\Omega(t)$, $v(t)$, $c(t)$, and $n_j(t)$ are unknown a priori and the latter three depend on $\bar{\lambda}(t)$, which is the value of the Lagrange multiplier vector used, according to Corollary 5.1, to compute the bids in the same time interval.

The objective is to adjust each element of $\bar{\lambda}(t)$ towards the smallest possible non-negative constant value for which the corresponding constraint among (6)-(8) is not violated. Define error signal $\bar{e}(t) = [e_1(t), e_{2,1}(t), \ldots, e_{2,m}(t), e_{3,2}(t), \ldots, e_{3,m}(t)]^T$, where

$$\begin{aligned} e_1(t) &= \xi_1/t_{max} - c(t), && (26) \\ e_{2,j}(t) &= \xi_{2,j} n_j(t) - c(t), \quad j = 1, \ldots, m, && (27) \\ e_{3,j}(t) &= n_j(t) - \xi_{3,j} n_1(t), \quad j = 2, \ldots, m. && (28) \end{aligned}$$

Constraints (6)-(8) are satisfied if and only if all elements of $\sum_t \bar{e}(t)$ are non-negative. Campaign controller achieves this by updating $\bar{\lambda}(t)$ based on constraint vector $\bar{\xi}$ and campaign feedback $\bar{y}(t) := [v(t), c(t), n_1(t), \cdots, n_m(t)]^T \in \mathbb{R}^{m+2}$. The value of $\bar{e}(t)$ for each individual time point is not relevant. In fact, sometimes it is important to allow $\bar{e}(t)$ to fluctuate over time in order for $\bar{\lambda}(t)$ to converge towards a constant. See e.g. [12] where feedforward control is used to distribute a daily spend budget throughout the day to make the convergence of the control signal and the Lagrange multiplier possible.

The plant is defined by the map $\bar{\lambda}(t) \mapsto \bar{e}(t)$; which in real applications is nonlinear, dynamic, time-varying, and stochastic [12]. Typically, the plant is approximately diagonal and linear in a neighborhood of each operating point. Refer to $\bar{a} := \text{diag}(\partial \bar{e}/\partial \bar{\lambda})$ evaluated at each operating point of $\bar{\lambda}$ as the plant gain vector, and assume a one-step plant delay.

B. Feedback Control

Consider pure I-error feedback control to update each element of $\bar{\lambda}(t)$ independently. Each Lagrange multiplier must be non-negative, hence the actuators are subject to saturation. Since pure integral control is used, integrator wind-up protection is easily obtained by using $\bar{\lambda}(t)$ as the state of the integral controller. That is, the feedback controllers are specified by

$$\lambda_1(t) = \max\big(\lambda_1(t-1) - c_{\text{I},1} e_1(t), 0\big), \quad (29)$$

$$\lambda_{2,j}(t) = \max\big(\lambda_{2,j}(t-1) - c_{\text{I},2,j} e_{2,j}(t), 0\big), \quad j = 1, \ldots, m, \quad (30)$$

$$\lambda_{3,j}(t) = \max\big(\lambda_{3,j}(t-1) - c_{\text{I},3,j} e_{3,j}(t), 0\big), \quad j = 2, \ldots, m, \quad (31)$$

where $c_{\text{I},1}, c_{\text{I},2,j}, c_{\text{I},3,j} \in \mathbb{R}_{>0}$ are the controller I *gains* of each feedback controller. Adequate values of these gains depend on the plant gains, which normally are estimated online. This estimation is outside the scope of this paper, but two approaches are discussed in [12] and summarized as:

Approach 1: Design an adaptive estimator, e.g. a recursive least squares algorithm, which based on time series observations $\bar{y}(t)$ estimates the elements of $\bar{a}$. This may require excitation control [12], which is a small random perturbation of $\bar{\lambda}(t)$ to ensure identifiability.

Approach 2: Use a large scale machine learning algorithm to process historical bid prices and context information offline and fit a prediction model of $F_{B_i^\star}(b)$ for the universe of all targeted impressions. Then use the techniques in section 'The Mathematics of Plant Smoothing' on pages 54-59 in [12] to compute an estimate of the plant gain.

Depending on how challenging the competitive landscape is for a specific campaign, it may be necessary to implement bid randomization (aka, Heisenberg bidding) [12] to make estimation and control possible.

Once $\bar{\lambda}(t)$ is updated, control signal vector $\bar{u}(t)$ is computed from $\bar{\lambda}(t)$ and $\bar{\xi}$ according to Corollary 5.1 and made available for real time bid calculation. Note, $\bar{\lambda}(t)$ and $\bar{u}(t)$ are computed at sparse discrete time points t while individual bids b_i^u are computed every time an impression opportunity is present, which may be billions of times per day.

VII. Simulation Example

The following example demonstrates the application of Theorem 5.1, Corollary 5.1, and the proof of concept campaign controller introduced in Section VI.

Suppose the campaign objective is to maximize the expected cumulative profit from ad-generated conversions (sold products), where each conversion is worth 10 dollars. Let C, N_1, N_2, and N_3 denote the spend and the number of impressions, conversions, and in-target impressions, respectively; where an in-target impression is defined by a specific context (based on e.g. website and Internet user features). The objective is subject to a spend constraint ξ_1, an effective cost per impression constraint $\xi_{2,1}$, an effective cost per conversion constraint $\xi_{2,2}$, and an in-target rate constraint $\xi_{3,3}$. All constraints are specified on a per day basis, and the sampling time of the implemented system is $\Delta = 5/60$ hours; i.e., $t_{max} = 24 \cdot 60/5 = 288$. Mathematically, the problem is

$$\underset{\{b_i \in \mathbb{R} | \forall i \in \Omega\}}{\text{maximize}} \quad 10\mathrm{E}N_2 - \mathrm{E}C$$

subject to

$$\mathrm{E}C \le \xi_1, \tag{32}$$
$$\mathrm{E}C \le \xi_{2,1}\mathrm{E}N_1, \tag{33}$$
$$\mathrm{E}C \le \xi_{2,2}\mathrm{E}N_2, \tag{34}$$
$$\xi_{3,3}\mathrm{E}N_1 \le \mathrm{E}N_3. \tag{35}$$

Suppose all impressions are traded in an auction exchange and sold according to a first price cost model.

The optimal bids $b_i, \forall i \in \Omega$, depend on the constraints $\xi_1, \xi_{2,1}, \xi_{2,2}$, and $\xi_{3,3}$; the event rates $p_{2,i}$ and $p_{3,i}$; the competitive landscape $F_{B_i^\star}(b)$; and the number of available impression opportunities (the cardinality of Ω). Assume impression valuation and bid shading optimization are perfect and consider how campaign controller handles changes in the constraints. All other aspects of the system are assumed fixed. The average spend constraint per sampling interval is denoted $\tilde{\xi}_1$, and is given by $\tilde{\xi}_1 = \xi_1/288$. Let the time-varying constraints be

$$\tilde{\xi}_1 = \begin{cases} 50, & t\Delta < 200 \\ 80, & 200 \le t\Delta < 800 \\ 25, & t\Delta \ge 800 \end{cases}, \qquad \xi_{2,1} = 4 \cdot 10^{-4},$$

$$\xi_{2,2} = \begin{cases} 5, & t\Delta < 600 \\ 3, & t\Delta \ge 600 \end{cases}, \qquad \xi_{3,3} = \begin{cases} 0.8, & t\Delta < 400 \\ 0.9, & t\Delta \ge 400 \end{cases}.$$

The campaign objective translates to an impression value $v_i = 10p_{2,i}$. Suppose $B_i^\star$ follows the Gamma(α_i, β_i) distribution, where α_i is the shape parameter and β_i is the inverse scale parameter [19]. A largest set of impression opportunities for which $v_i, p_{2,i}, p_{3,i}$, and $F_{B_i^\star}(b)$ are fixed defines a context. Assume, for simplicity, that there are only four contexts (types) of impressions. Furthermore, suppose the number of impression opportunities per time interval and context is constant. Let the contexts be characterized by

Context of i	$p_{2,i}$	$p_{3,i}$	α_i	β_i	card(Ω_i)
1	$1 \cdot 10^{-4}$	0.2	4.9	$3.3 \cdot 10^3$	$3 \cdot 10^5$
2	$5 \cdot 10^{-5}$	0.5	44	$4.4 \cdot 10^4$	$2 \cdot 10^5$
3	$2 \cdot 10^{-5}$	0.8	44	$5.9 \cdot 10^4$	$1 \cdot 10^5$
4	$1 \cdot 10^{-4}$	0.9	44	$1.8 \cdot 10^5$	$1.5 \cdot 10^5$

Figure 3 shows visually the shape of the CDF $F_{B_i^\star}(b)$ for each different context. From the shape of the CDFs, it is clear that the relative standard deviation of the highest competing bid is much larger for context 1 impressions than it is for other impression opportunities.

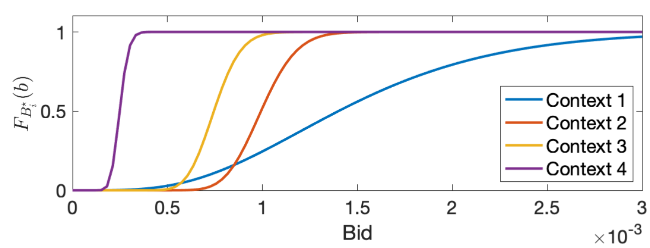

Fig. 3. The CDF $F_{B_i^\star}(b)$ of the highest competing bid price. Note, $F_{B_i^\star}(b)$ depends on the context of the ith impression opportunity.

The responsibility of the controller is to update $\lambda_1(t)$, $\lambda_{2,1}(t)$, $\lambda_{2,2}(t)$, and $\lambda_{3,3}(t)$ based on the error feedback signals defined by (26)-(28). These updates are governed by four separate integral controllers with wind-up protection as specified by (29)-(31), where the I-gains are chosen largely ad hoc to be $c_{\text{I},1} = 0.1/\tilde{\xi}_1$, $c_{\text{I},2,j} = 0.3/\xi_{2,j}$, $j = 1, 2$, and $c_{\text{I},3,3} = 10^{-8}/\xi_{3,3}$. In a practical implementation, the I-gains are selected based on our best understanding of the plant gains, which may be computed according to Approaches 1 or 2 discussed in the previous section. As for the error feedback, $c(t)$ is the observed spend and $n_j(t)$, for $j = 1, 2, 3$, are the observed event counts in the most recent time interval.

Figure 4 shows the result of simulating the above closed loop system for 1000 hours when $\bar{\lambda}(0)$ is initialized as the zero vector and there is no un-modeled dynamics and no noise. Subplot 1,1 displays the bid price for each impression type, subplot 1,2 shows the overall win rate of impression opportunities, and subplot 1,3 shows the total ad value. The

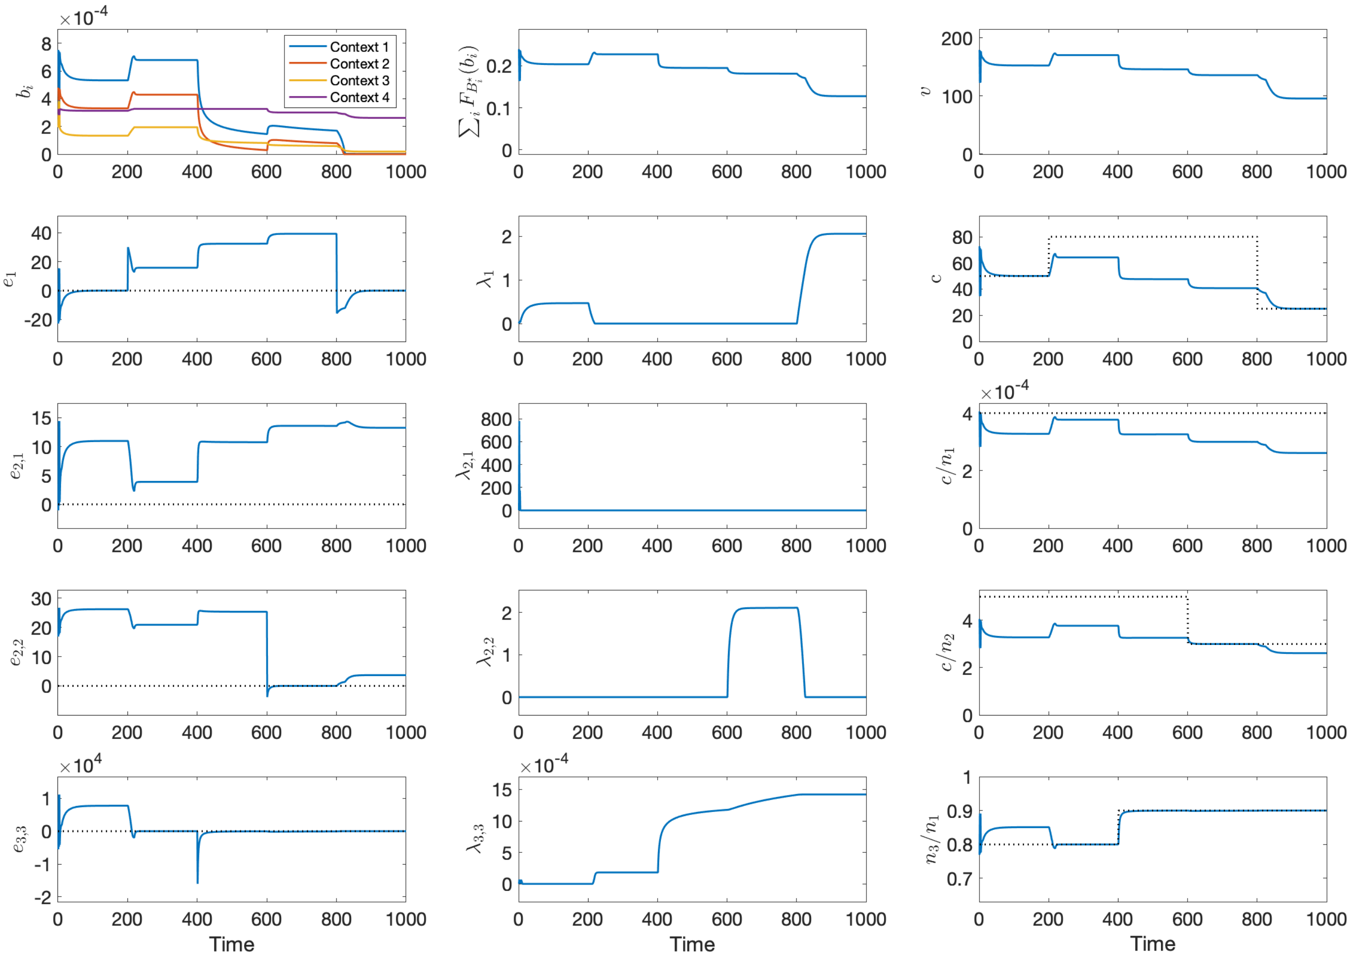

Fig. 4. Closed-loop performance of the multi-objective optimization problem using four separate pure I-feedback controllers in a noise-free environment.

other subplots show how errors (26)-(28) and Lagrange multipliers (29)-(31) evolve over time; and how $c(t)$, $c(t)/n_1(t)$, $c(t)/n_2(t)$, and $n_3(t)/n_1(t)$ compare to $\tilde{\xi}_1$, $\xi_{2,1}$, $\xi_{2,2}$, and $\xi_{3,3}$, respectively.

Initially, the spend rate $c(t)$ and the average cost per impression $c(t)/n_1(t)$ are prohibitively large, while the in-target rate $n_3(t)/n_1(t)$ is too small. This triggers the controller to adjust $\lambda_1(t)$, $\lambda_{2,1}(t)$, and $\lambda_{3,3}(t)$ to strictly positive values. However, after a short time and several adjustments, all Lagrange multipliers settle on constant values where only the spend constraint $\tilde{\xi}_1$ is restricting the delivery. I.e., only $\lambda_1(t)$ is larger than zero. At time $t\Delta = 200$, budget $\tilde{\xi}_1$ is increased, but this larger budget cannot be fully spent due to the in-target event rate constraint $\xi_{3,3}n_1(t) \leq n_3(t)$, which becomes the only binding constraint. The Lagrange multipliers settle on a new configuration of constant values where only $\lambda_{3,3}(t)$ is larger than zero. The in-target rate constraint $\xi_{3,3}$ is increased at time $t\Delta = 400$ causing bids for context 1, 2, and 3 impressions to drop to a level where most of the spend is on context 4 impressions with an in-target rate $p_{3,i} = 0.9$. The cost per conversion constraint $\xi_{2,2}$ is reduced at time $t\Delta = 600$ leading to $c(t) = \xi_{2,2}n_2(t)$ and $\xi_{3,3}n_1(t) = n_3(t)$ being simultaneously the only binding constraints. Finally, at $t\Delta = 800$ the budget is reduced dramatically leading to the spend and in-target rate constraints being the only two binding constraints.

Based on what constraints are binding, it follows from Theorem 5.1 that the optimal adjusted impression value b_i^u in the different time intervals takes the form

$$b_i^u = \begin{cases} \dfrac{10p_{2,i}}{1+\lambda_1}, & t\Delta < 200, \\ 10p_{2,i} + (p_{3,i} - \xi_{3,3})\lambda_{3,3}, & t\Delta \in [200, 600), \\ \dfrac{10p_{2,i} + p_{2,i}\xi_{2,2}\lambda_{2,2} + (p_{3,i} - \xi_{3,3})\lambda_{3,3}}{1+\lambda_{2,2}}, & t\Delta \in [600, 800), \\ \dfrac{10p_{2,i} + (p_{3,i} - \xi_{3,3})\lambda_{3,3}}{1+\lambda_1}, & t\Delta \geq 800. \end{cases}$$

A simpler interface between impression valuation and campaign control (Figure 2) is obtained by invoking Corollary 5.1, which states that $b_i^u = [1, \bar{v}_i^T]\bar{u}$, where $\bar{u} \in \mathbb{R}^{m+1}$ is a rational function of $\bar{\xi}$ and $\bar{\lambda}$. The elements of $\bar{u}(t)$ are plotted in Figure 5. Note that $u_0(t)$ in this example is negative. This illustrates that for some impression opportunities, in a real application, $b_i^u < 0$ which is mapped to $b_i^{opt} = 0$. However, $B_i^\star > 0$, and no impression is awarded for a bid that is not strictly positive.

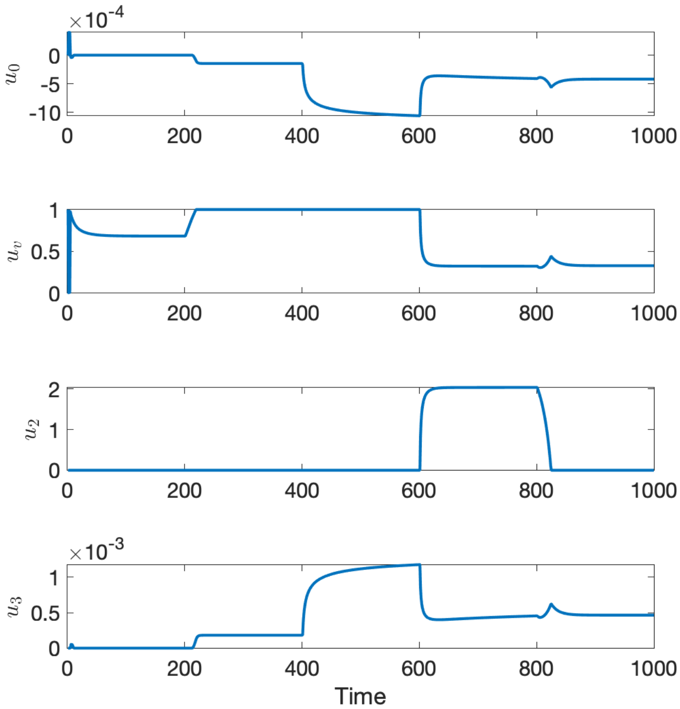

Fig. 5. The elements of $\bar{u}(t) = [u_0, u_v, u_2, u_3]^T$ versus time.

VIII. Conclusions and Future Work

Lagrangian techniques together with probabilistic tools were used to derive the optimal bidding mechanism to solve a multi-objective optimization problem in programmatic advertising. It is demonstrated how the bid calculation lends itself to a decoupled implementation involving three separate and scalable systems dealing with impression valuation, campaign control, and bid shading optimization.

The control system plays a critical role in the implementation of the optimization system, and enables a scalable and optimal solution. As a proof of concept, the multiple-input, multiple-output control problem is implemented as separate single-input, single-output pure I-error feedback controllers and evaluated in a noise-free simulated scenario. The closed-loop performance is according to intuition and theory and serves as a basic benchmark test case.

The plan for the future is to first of all research algorithms for robust system identification. Being able to estimate plant gains, dynamics, time-varying effects, and noise characteristics is important before the solution can be used in a fully automated setting. With solid algorithms for system identification in place, the next step is to refine the control design.

Theorem 5.1 provides necessary conditions for optimality, but the objective function and the constraints are not in general convex, hence there may exist other bids satisfying the conditions. An important future research problem is to establish useful conditions that are both necessary and sufficient for optimality, or bounds on how much the expected surplus is compromised if a suboptimal solution is used.

Another relevant research problem is to relax the assumptions that $N_{j,i}$ and $B_i^\star$, as well as, V_i and $B_i^\star$ are conditionally independent, given the context. The assumption is reasonable in practice at most if the context definition is granular (rich in features). As privacy concerns are growing across society, policies and laws are implemented to reduce what is known about each Internet user and each impression opportunity. For this reason, the above independence assumption is not likely to hold true in the future.

References

[1] Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, pages 1307–1315, New York, NY, USA, 2011. ACM.

[2] Grégoire Jauvion and Nicolas Grislain. Optimal allocation of real-time-bidding and direct campaigns. In *Proceedings of the 24th ACM SIGKDD Int. Conference on Knowledge Discovery & Data Mining*, KDD '18, pages 416–424, New York, NY, USA, 2018. ACM.

[3] Weinan Zhang, Shuai Yuan, and Jun Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pages 1077–1086, New York, NY, USA, 2014. ACM.

[4] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. Bid optimization by multivariable control in display advertising. *CoRR*, abs/1905.10928, 2019.

[5] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. Deep & cross network for ad click predictions. *CoRR*, abs/1708.05123, 2017.

[6] Junwei Pan, Jian Xu, Alfonso Lobos Ruiz, Wenliang Zhao, Shengjun Pan, Yu Sun, and Quan Lu. Field-weighted factorization machines for click-through rate prediction in display advertising. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pages 1349–1357, 2018.

[7] Yuya Yoshikawa and Yusaku Imai. A nonparametric delayed feedback model for conversion rate prediction. *CoRR*, abs/1802.00255, 2018.

[8] Niklas Karlsson and Jianlong Zhang. Applications of feedback control in online advertising. In *Proceedings of the 2013 American Control Conference, Washington, DC, June 17 - 19*, pages 6008 – 6013, 2013.

[9] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. Budget pacing for targeted online advertisements at linkedin. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pages 1613–1619, New York, NY, USA, 2014. ACM.

[10] Niklas Karlsson. Control of periodic systems in online advertising. In *Proceedings of the 2018 IEEE 57th Conference on Decision and Control, Miami, FL, USA*, Dec 17-19 2018.

[11] Niklas Karlsson. Adaptive optimization and control in online advertising. In *Proceedings of the 2019 IEEE 58th Conference on Decision and Control, Nice, France*, Dec 11-13 2019.

[12] Niklas Karlsson. Feedback control in programmatic advertising: The frontier of optimization in real-time bidding. *IEEE Control Systems Magazine*, 40(5):40–77, 2020.

[13] Djordje Gligorijevic, Tian Zhou, Bharatbhushan Shetty, Brendan Kitts, Shengjun Pan, Junwei Pan, and Aaron Flores. Bid shading in the brave new world of first-price auctions. *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, Oct 2020.

[14] Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhushan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin, and Jianlong Zhang. Bid shading by win-rate estimation and surplus maximization. 2020.

[15] Niklas Karlsson and Qian Sang. Adaptive bid shading optimization of first-price ad inventory,. In *Proceedings of the 2021 American Control Conference, New Orleans, LA, USA*, ACC '21, pages 4971–4978, May 26–28 2021.

[16] Tian Zhou, Hao He, Shengjun Pan, Niklas Karlsson, Bharatbhushan Shetty, Brendan Kitts, Djordje Gligorijevic, San Gultekin, Tingyu Mao, Junwei Pan, Jianlong Zhang, and Aaron Flores. An efficient deep distribution network for bid shading in first-price auctions. In *In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'21), August 14-18, 2021, Singapore*, 2021.

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

[18] David G. Luenberger and Yinyu Ye. *Linear and Nonlinear Programming*. Springer Publishing Company, Incorporated, 2015.

[19] George Casella and Roger L. Berger. *Statistical Inference*. Duxbury, 2 edition, 2001.