2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Feedback-control based Hierarchical Multi-constraint Ad Campaign Optimization

Niklas Karlsson[1]

***Abstract*— Online advertising is typically implemented via real-time bidding, and advertising campaigns are then defined as extremely high-dimensional optimization problems. Advertisers often define a campaign by an *order* consisting of multiple *lines*. Campaign delivery constraints may be imposed on the order as a whole and on each ad line. E.g., there may be budget and cost per click constraints on the order and on each line individually. Furthermore, the sum of line budgets may exceed the order budget, and the cost per click constraint on lines may differ. This leaves room for cross-line budget optimization; i.e., budget may be shifted across lines to maximize the advertising value without violating the constraints. This paper derives the optimal bidding mechanism for a large family of constrained optimization problems. It is shown how the optimal bidding strategy can be implemented as scalable non-cooperating agents on the order and the individual lines.**

***Index Terms* — Optimization, Real-time Bidding, Programmatic Advertising**

I. INTRODUCTION

The business model of many Internet companies is centered around online advertising, where ad *impressions*, which are views of ads, are traded on open exchanges. A *Demand Side Platform* (DSP) is a particular business model that serves as the middleman between an advertiser and impression exchanges. The DSP provides bid optimization as a service and helps advertisers to spend their advertising budgets optimally. A scalable implementation of this optimization involves feedback control as a key component.

It is common that an advertiser strives to reach multiple audiences (sometimes with different ad creatives), but with a finite total budget. Managing the advertising objective for each audience is handled by introducing the concepts of *order* and *line*, where an order consists of multiple lines. An order constraint dictates a constraint that must be satisfied across lines (e.g. a total budget), while line constraints prescribe constraints pertaining only to a line (e.g. a line budget). The advertiser would like to maximize the total advertising profit or value of the campaign, but with the caveat that delivery constraints may be imposed on the order or on the individual lines. Historically, this allocation of budget to lines has been handled manually by experienced account managers, or by some non-feedback based planner. However, these approaches are typically suboptimal since they do not efficiently learn from mistakes.

There is a rich literature in constrained optimization of display advertising. Simple constraints with unrealistic assumptions are treated in [1], [2]. Bid optimization involves estimating e.g. *click-through rate* (CTR) and *conversion rate* (CVR), e.g. [3], [4], [5] and many others; as well as feedback control-based bid adjustments to ensure various ad campaign constraints are satisfied [6], [7], [8], [9], [10]. Whenever the impressions are sold based on a first price cost model, the optimization also involves bid shading [11], [12], [13], [14].

The most similar work to the current paper is [15], which deals with non-hierarchical advertiser specified optimization constraints, and with a variation published in [16] to incorporate a DSP specified cost per bid constraint. Our contribution is the generalization of the result in [15] to support campaign constraints configured hierarchically, and the novelty lies in the modularization of the bidding strategy that enables the solution being implemented as non-cooperating agents on the order and the individual lines.

The article is organized as follows. Section II introduces key notation while Section III defines the problem that is solved later and Section IV introduces some additional notation. The key contribution and the optimal bidding mechanism is presented in Section V. Section VI describes how the theoretical results justify an architecture involving four non-cooperating subsystems. The architecture is made clear with help of an example in Section VII, and conclusions and ideas of future work are discussed in Section VIII.

II. SETUP AND NOTATION

The optimization problem considered in this paper is for a specific ad campaign, and the objective is to derive the optimal strategy for real time bidding on impressions. 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 for each impression opportunity is decided by the seller of the impression and is known to the bid engine (the DSP) before a bid is computed and submitted to the auction. 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. Key notation used in the paper is:

- The campaign is defined by an *order*, which is composed of *lines* $\ell = 1, \ldots, \ell_{max}$.
- The *bid price* $b_{\ell,i} \in [0, \infty)$, $i \in \Omega^\ell$, $\ell = 1, \ldots, \ell_{max}$, is the bid amount submitted by line ℓ for the ith impression opportunity, and the set of all $b_{\ell,i}$ are decision variables of the optimization problem.
- The set of impression opportunities that line ℓ may bid on is denoted Ω^ℓ, whereas the set of all impression opportunities is denoted Ω^0; i.e., $\Omega^0 = \bigcup_{\ell=1}^{\ell_{max}} \Omega^\ell$. Suppose $\Omega^\ell \cap \Omega^k = \emptyset$, if $0 \neq \ell \neq k \neq 0$, which means no two lines within the campaign compete over the same impressions. Furthermore, assume Ω^ℓ may be partitioned into subsets Ω_1^ℓ and Ω_2^ℓ denoting impression opportunities sold based on a first and second price cost model, respectively. Let $\Omega_1^\ell \cap \Omega_2^\ell = \emptyset$ and $\Omega_1^\ell \cup \Omega_2^\ell = \Omega^\ell$. In practice, some web publishers sell all their ad impressions based on a first price cost model ($i \in \Omega_1^\ell$), while others sell all impressions based on a second price

[1]N. Karlsson (nkarlsson8 at gmail.com) is Senior Principal Research Scientist with Amazon Ads, Palo Alto, California, USA.

cost model ($i \in \Omega_2^\ell$). The cost model for a specific impression is known ahead of time. However, the cardinality of Ω_1^ℓ and Ω_2^ℓ are not a priori known since they reflect the future Internet traffic (the aggregate number web site visits by Internet users).

- The uppercase of a Roman letter denotes a random variable, while a lowercase represents its expected value (future event) or realization (historical event).
- 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. Impression $i \in \Omega^\ell$ is awarded to line ℓ if $b_{\ell,i} \geq b_i^\star$.
- 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 of [15] for a few examples. Note, $F_{B_i^\star}(b)$ is typically not known ahead of time, but must be predicted based on historical data. Exactly what data can be used for this prediction depends on what auction data is shared by the exchange after the auction for an impression is completed.
- The *ad cost*, $C_{\ell,i}$, of an awarded impression $i \in \Omega^\ell$ is the amount ad line ℓ is charged for the impression. It follows from the definition of the two cost models that $C_{\ell,i} = b_{\ell,i}$, if $i \in \Omega_1^\ell$, and $C_{\ell,i} = B_i^\star$, if $i \in \Omega_2^\ell$.
- The *event count* $N_{\ell,j,i} \in \mathbb{N}$ is the number of events of type $j \in \{1, \ldots, j_{max}\}$ following an impression i awarded to line ℓ, where j_{max} is the number of event types under consideration. Examples of event types are impressions, clicks, and conversions (product sales).
- The *impression value* $V_{\ell,i} \in \mathbb{R}_{>0}$ is the advertising value attributed to the ith impression, if awarded. It encodes branding and/or performance value, and is typically a function of one or more of the event counts.
- The line-level *total number of j-events*, $N_{\ell,j}$, the *total value*, V_ℓ, and the *cumulative ad cost*, C_ℓ, are

$$N_{\ell,j} = \sum_{i \in \Omega^\ell} N_{\ell,j,i} \mathbb{I}_{\{b_{\ell,i} \geq B_i^\star\}}, \qquad j = 1, \ldots, j_{max} \quad (1)$$

$$V_\ell = \sum_{i \in \Omega^\ell} V_{\ell,i} \mathbb{I}_{\{b_{\ell,i} \geq B_i^\star\}}, \quad (2)$$

$$C_\ell = \sum_{i \in \Omega^\ell} C_{\ell,i} \mathbb{I}_{\{b_{\ell,i} \geq B_i^\star\}}; \quad (3)$$

where $\mathbb{I}_A$ is the indicator function satisfying $\mathbb{I}_A = 1$, if $A =$ true, and $\mathbb{I}_A = 0$, otherwise, and the campaign (aka, order-level) total event count, value, and cost are

$$N_{0,j} = \sum_{\ell=1}^{\ell_{max}} N_{\ell,j}, \quad (4)$$

$$V_0 = \sum_{\ell=1}^{\ell_{max}} V_\ell, \quad (5)$$

$$C_0 = \sum_{\ell=1}^{\ell_{max}} C_\ell. \quad (6)$$

- The expected impression value and event count, given the context, are $\mathrm{E}V_{\ell,i} = v_{\ell,i}$ and $\mathrm{E}N_{\ell,j,i} = p_{\ell,j,i}$ (breaking from the convention that suggests the notation $n_{\ell,j,i}$). Typically, $N_{\ell,j,i}$ is a binary random variable, which makes $p_{\ell,j,i}$ a probability, or success rate, between zero and one. Event type $j = 1$ is without loss of generality an impression, which implies that $p_{\ell,1,i} = 1$, for $\ell = 1, \ldots, \ell_{max}$ and for all $i \in \Omega$.
- The random variables $N_{\ell,j,i}$ and $B_i^\star$, as well as, the random variables $V_{\ell,i}$ and $B_i^\star$ are assumed conditionally independent, given the context.

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

III. Problem Formulation

The objective is to solve an advertiser-specified optimization problem, e.g. to maximize the return on investment on an advertising campaign subject to a fixed budget. A general such problem is to maximize the expected campaign-aggregate total cost-discounted profit $\mathrm{E}V_0 - \alpha \mathrm{E}C_0$, for *cost discount* parameter $\alpha \geq 0$. Parameter α is defined by the advertiser based on their preference towards value or profit. The campaign is defined by an order consisting of ℓ_{max} ad lines. The order ($\ell = 0$) as a whole and each line ($\ell = 1, \ldots, \ell_{max}$) separately are subject to constraints on total ad cost (spend), *effective cost per event* (eCPX), *events per impression rate* (ER), specified by $\xi_{\ell,1}, \xi_{\ell,2,j}, \xi_{\ell,3,j} \geq 0$; and a max bid constraint, $b_{max} > 0$, also referred to as "max bid," where j denotes the event type (e.g. click or conversion). The problem is mathematically defined by

$$\underset{\{b_{\ell,i} | 0 \leq b_{\ell,i} \leq b_{max}, \forall i \in \Omega^\ell, \ell = 1, \ldots, \ell_{max}\}}{\text{maximize}} \quad \mathrm{E}V_0 - \alpha \mathrm{E}C_0 \quad (7)$$

subject to

$$\mathrm{E}C_\ell \leq \xi_{\ell,1} \qquad \text{(spend)} \quad (8)$$

$$\mathrm{E}C_\ell \leq \xi_{\ell,2,j} \mathrm{E}N_{\ell,j} \qquad \text{(eCPX)} \quad (9)$$

$$\xi_{\ell,3,j} \mathrm{E}N_{\ell,1} \leq \mathrm{E}N_{\ell,j} \qquad (\text{ER}, j \neq 1) \quad (10)$$

where the constraints are defined for $\ell = 0, 1, \ldots, \ell_{max}$ and $j = 1, 2, \ldots, j_{max}$. Note, $j = 1$, by convention, corresponds to an impression event. Hence, $N_{\ell,1}$ represents the total number of awarded impressions, which renders (10) meaningless for $j = 1$. The cardinality of $\Omega^0 = \bigcup_{\ell=1}^{\ell_{max}} \Omega^\ell$ is in the order of millions or billions making the problem extremely high-dimensional. Moreover, the cardinality of Ω^0, the impression value $v_{\ell,i}$, the event rate $p_{\ell,j,i}$, and the minimum bid to win $b_i^\star$ are a priori unknown. This makes it impractical to solve (7)-(10) using a centralized method.

IV. Preliminaries

Define *constraint vector*
$\bar{\xi}_\ell = [\xi_{\ell,1}, \xi_{\ell,2,1}, \cdots, \xi_{\ell,2,j_{max}}, \xi_{\ell,3,2}, \cdots, \xi_{\ell,3,j_{max}}]^T$,
Lagrange multiplier vector
$\bar{\lambda}_\ell = [\lambda_{\ell,1}, \lambda_{\ell,2,1}, \cdots, \lambda_{\ell,2,j_{max}}, \lambda_{\ell,3,2}, \cdots, \lambda_{\ell,3,j_{max}}]^T$,
and *impression value vector*
$\bar{v}_{\ell,i} = [v_{\ell,i}, p_{\ell,2,i}, p_{\ell,3,i}, \cdots, p_{\ell,j_{max},i}]^T$ where $i \in \Omega^\ell$ and $\ell = 0, \ldots, \ell_{max}$. Vector $\bar{\xi}_\ell$ contains order/line-level constraints; vector $\bar{\lambda}_\ell$ is composed of multipliers, which in an implementation are used as tuning knobs and adjusted gracefully over time toward their optimal values; and vector

$\bar{v}_{\ell,i}$ consists of the expected impression value and event rates (e.g. click-through rate). Furthermore, define helper functions

$$\begin{aligned} g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell) &= v_{\ell,i} + \sum_{j=1}^{j_{max}} p_{\ell,j,i}\xi_{\ell,2,j}\lambda_{\ell,2,j} \\ &\quad + \sum_{j=2}^{j_{max}} (p_{\ell,j,i} - \xi_{\ell,3,j})\lambda_{\ell,3,j}, \\ g_1(\bar{\lambda}_\ell) &= \alpha + \lambda_{\ell,1} + \sum_{j=1}^{j_{max}} \lambda_{\ell,2,j}. \end{aligned}$$

The above notation is next used in the proof of the main result of the paper, which are conditions for optimal bidding.

V. Optimal Bidding Mechanism

The following theorem states necessary conditions for optimal bidding, and justifies a decomposition of the bid calculation into a solution that involves $1 + \ell_{max}$ non-cooperating feedback controllers.

Theorem 5.1: The optimal bid prices $b^{opt}_{\ell,i}$, for all $i \in \Omega^\ell$ and $\ell = 1, \ldots, \ell_{max}$, satisfy $b^{opt}_{\ell,i} = b_{max}\mathbb{I}_{\{g_0(\bar{v}_{\ell,i},\bar{\xi}_\ell,\bar{\lambda}_\ell)\geq 0\}}$, if $g_1(\bar{\lambda}_\ell) = 0$, and

$$b^{opt}_{\ell,i} = \underset{0\leq b\leq b_{max}}{\operatorname{argmax}} \ \mathrm{E}\tilde{S}_{\ell,i}(b; b^u_{\ell,i}); \qquad (11)$$

otherwise, where *surplus per bid response*, $\tilde{S}_{\ell,i}$, and *adjusted impression value*, $b^u_{\ell,i}$, are defined by

$$\tilde{S}_{\ell,i}(b; b^u_{\ell,i}) = \left(b^u_{\ell,i} - C_{\ell,i}(b)\right)\mathbb{I}_{\{b\geq B^\star_i\}}, \qquad (12)$$

$$b^u_{\ell,i} = \frac{g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell)}{g_1(\bar{\lambda}_\ell)}; \qquad (13)$$

where, for all values of ℓ and j, inequalities (8)-(10) hold, as well as, $\lambda_{\ell,1} \geq \lambda_{0,1} \geq 0$, $\lambda_{\ell,2,j} \geq \lambda_{0,2,j} \geq 0$, $\lambda_{\ell,3,j} \geq \lambda_{0,3,j} \geq 0$; and $\lambda_{0,1}(\mathrm{E}C_0 - \xi_0) = (\lambda_{\ell,1} - \lambda_{0,1})(\mathrm{E}C_\ell - \xi_\ell) = \lambda_{0,2,j}(\mathrm{E}C_0 - \xi_{0,2,j}\mathrm{E}N_{0,j}) = (\lambda_{\ell,2,j} - \lambda_{0,2,j})(\mathrm{E}C_\ell - \xi_{\ell,2,j}\mathrm{E}N_{\ell,j}) = \lambda_{0,3,j}(\xi_{0,3,j}\mathrm{E}N_{0,1} - \mathrm{E}N_{0,j}) = (\lambda_{\ell,3,j} - \lambda_{0,3,j})(\xi_{\ell,3,j}\mathrm{E}N_{\ell,1} - \mathrm{E}N_{\ell,j}) = 0$.

Proof: The Lagrangian of (7)-(10) is $\mathcal{L} = \mathrm{E}V_0 - \alpha\mathrm{E}C_0 - \sum_{\ell=0}^{\ell_{max}}\tilde{\lambda}_{\ell,1}(\mathrm{E}C_\ell - \xi_{\ell,1}) - \sum_{\ell=0}^{\ell_{max}}\sum_{j=1}^{j_{max}}\tilde{\lambda}_{\ell,2,1}(\mathrm{E}C_\ell - \xi_{\ell,2,j}\mathrm{E}N_{\ell,j}) - \sum_{\ell=0}^{\ell_{max}}\sum_{j=2}^{j_{max}}\tilde{\lambda}_{\ell,3,1}(\xi_{\ell,3,j}\mathrm{E}N_{\ell,1} - \mathrm{E}N_{\ell,j})$. If there exist $b_{\ell,i}$, $\tilde{\lambda}_{\ell,1}, \tilde{\lambda}_{\ell,2,j}, \tilde{\lambda}_{\ell,3,j} \geq 0$, $\forall i \in \Omega^\ell$, $\ell = 0, \ldots, \ell_{max}$, and $j = 1, 2, \ldots, j_{max}$, such that the $b_{\ell,i}$'s maximize $\mathcal{L}$, $b_{\ell,i} \leq b_{max}$, the inequalities (8)-(10) are satisfied, and $\tilde{\lambda}_{\ell,1}(\mathrm{E}C_\ell - \xi_\ell) = \tilde{\lambda}_{\ell,2,j}(\mathrm{E}C_\ell - \xi_{\ell,2,j}\mathrm{E}N_{\ell,j}) = \tilde{\lambda}_{\ell,3,j}(\xi_{\ell,3,j}\mathrm{E}N_{\ell,1} - \mathrm{E}N_{\ell,j}) = 0$; then, due to the Lagrangian sufficiency theorem [18], these values of $b_{\ell,i}$ solve (7)-(10). Use relationships (4)-(6), and rearrange the terms to obtain $\mathcal{L} = \mathrm{E}(\sum_{\ell=0}^{\ell_{max}}\tilde{\lambda}_{\ell,1}\xi_{\ell,1} + \sum_{\ell=1}^{\ell_{max}} V_\ell - \alpha\sum_{\ell=1}^{\ell_{max}} C_\ell - \sum_{\ell=1}^{\ell_{max}}(\tilde{\lambda}_{0,1} + \tilde{\lambda}_{\ell,1})C_\ell - \sum_{\ell=1}^{\ell_{max}}\sum_{j=1}^{j_{max}}(\tilde{\lambda}_{0,2,j} + \tilde{\lambda}_{\ell,2,j})(C_\ell - \xi_{\ell,2,j}N_{\ell,j}) - \sum_{\ell=1}^{\ell_{max}}\sum_{j=2}^{j_{max}}(\tilde{\lambda}_{0,3,j} + \tilde{\lambda}_{\ell,3,j})(\xi_{\ell,3,j}N_{\ell,1} - N_{\ell,j}))$.

Define $\lambda_{0,1} = \tilde{\lambda}_{0,1}$, $\lambda_{\ell,1} = \tilde{\lambda}_{0,1} + \tilde{\lambda}_{\ell,1}$, $\lambda_{\ell,2,j} = \tilde{\lambda}_{0,2,j} + \tilde{\lambda}_{\ell,2,j}$, and $\lambda_{\ell,3,j} = \tilde{\lambda}_{0,3,j} + \tilde{\lambda}_{\ell,3,j}$; and collect all terms of V_ℓ, C_ℓ, $N_{\ell,j}$. It follows that $\mathcal{L} = \sum_{\ell=0}^{\ell_{max}}(\lambda_{\ell,1} - \lambda_{0,1})\xi_{\ell,1} + \sum_{\ell=1}^{\ell_{max}}\mathrm{E}(V_\ell + \sum_{j=1}^{j_{max}}\lambda_{\ell,2,j}\xi_{\ell,2,j}N_{\ell,j} + \sum_{j=2}^{j_{max}}\lambda_{\ell,3,j}(N_{\ell,j} - \xi_{\ell,3,j}N_{\ell,1}) - (\alpha + \lambda_{\ell,1} + \sum_{j=1}^{j_{max}}\lambda_{\ell,2,j})C_\ell)$.

Substitute for $N_{\ell,j}, V_\ell$, and C_ℓ, as given in (1)-(3), and swap the order of summation over i and expectation.

$\mathcal{L} = \sum_{\ell=0}^{\ell_{max}}(\lambda_{\ell,1} - \lambda_{0,1})\xi_{\ell,1} + \sum_{\ell=1}^{\ell_{max}}\sum_{i\in\Omega^\ell}\mathrm{E}(V_{\ell,i} + \sum_{j=1}^{j_{max}}\lambda_{\ell,2,j}\xi_{\ell,2,j}N_{\ell,1,i} + \sum_{j=2}^{j_{max}}\lambda_{\ell,3,j}(N_{\ell,j,i} - \xi_{\ell,3,j}N_{\ell,1,i}) - (\alpha + \lambda_{\ell,1} + \sum_{j=1}^{j_{max}}\lambda_{\ell,2,j})C_{\ell,i})\mathbb{I}_{\{b_{\ell,i}\geq B^\star_i\}}$.

Random variables $N_{\ell,j,i}$ and $V_{\ell,i}$ are by assumption conditionally independent of $B^\star_i$, given the context (see Section II), which implies $\mathrm{E}(N_{\ell,j,i}\mathbb{I}_{\{b_i\geq B^\star_i\}}) = \mathrm{E}(N_{\ell,j,i})\mathrm{E}(\mathbb{I}_{\{b_i\geq B^\star_i\}}) = p_{\ell,j,i}\mathrm{E}(\mathbb{I}_{\{b_i\geq B^\star_i\}})$ and $\mathrm{E}(V_{\ell,i}\mathbb{I}_{\{b_i\geq B^\star_i\}}) = \mathrm{E}(V_{\ell,i})\mathrm{E}(\mathbb{I}_{\{b_i\geq B^\star_i\}}) = v_{\ell,i}\mathrm{E}(\mathbb{I}_{\{b_i\geq B^\star_i\}})$. Moreover, $p_{\ell,1,i} \equiv 1$. Using helper functions $g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell)$ and $g_1(\bar{\lambda}_\ell)$ introduced in Section IV, it follows after a simple rearrangement that $\mathcal{L} = \sum_{\ell=0}^{\ell_{max}}(\lambda_{\ell,1} - \lambda_{0,1})\xi_{\ell,1} + \sum_{\ell=1}^{\ell_{max}}\sum_{i\in\Omega^\ell}\mathrm{E}(g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell) - g_1(\bar{\lambda}_\ell)C_{\ell,i})\mathbb{I}_{\{b_{\ell,i}\geq B^\star_i\}}$.

For fixed values of $\lambda_{\ell,1}, \lambda_{\ell,2,j}$, and $\lambda_{\ell,3,j}$, the Lagrangian $\mathcal{L}$ may be optimized for each i independently by finding the $b_i \in [0, b_{max}]$ that maximizes $\mathcal{L}_{\ell,i}$, where $\mathcal{L}_{\ell,i} = \mathrm{E}(g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell) - g_1(\bar{\lambda}_\ell)C_{\ell,i})\mathbb{I}_{\{b_{\ell,i}\geq B^\star_i\}}$.

Next, consider $g_1(\bar{\lambda}_\ell) = 0$ and $g_1(\bar{\lambda}_\ell) > 0$ separately.

Case 1: If $g_1(\bar{\lambda}_\ell) = 0$, then $\mathcal{L}_i = \mathrm{E}\big(g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell)\mathbb{I}_{\{b_{\ell,i}\geq B^\star_i\}}\big) = g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell)F_{B^\star_i}(b_i)$, which is a non-decreasing function of $b_{\ell,i}$ whenever $g_0(\bar{v}_{\ell,i}, \bar{\xi}_\ell, \bar{\lambda}_\ell) \geq 0$, and is a non-increasing function otherwise. Therefore, $b^{opt}_{\ell,i} = b_{max}\mathbb{I}_{\{g_0(\bar{v}_{\ell,i},\bar{\xi}_\ell,\bar{\lambda}_\ell)\geq 0\}}$.

Case 2: If $g_1(\bar{\lambda}_\ell) > 0$, then $\mathcal{L}_i = g_1(\bar{\lambda}_\ell)\mathrm{E}\big((b^u_{\ell,i} - C_{\ell,i}(b_{\ell,i}))\mathbb{I}_{\{b_{\ell,i}\geq B^\star_i\}}\big) = g_1(\bar{\lambda}_\ell)\mathrm{E}\tilde{S}_{\ell,i}(b_{\ell,i}; b^u_{\ell,i})$, where $b^u_{\ell,i}$ and $\tilde{S}_{\ell,i}$ are defined in (13) and (12). Finally, note that $\mathcal{L}_{\ell,i}$ and $\mathrm{E}\tilde{S}_{\ell,i}$ are maximized for the same value of b, hence, $b^{opt}_{\ell,i} = \underset{0\leq b\leq b_{max}}{\operatorname{argmax}}\ \mathrm{E}\tilde{S}_{\ell,i}(b; b^u_{\ell,i})$, which completes the proof. ■

Theorem 5.1 justifies how problem (7)-(10) can be decomposed and solved as a multi-agent non-cooperative game, where each agent solves a much simpler problem.

VI. Application

An ad campaign has a budget, a flight time, and a goal that encodes performance and/or branding objectives. The goal is defined by a profit function and by delivery constraints that are related to e.g. spend, cost per click, or the ratio of impressions served to users from a particular demographic. The objective is to bid optimally on impression opportunities that become available throughout the flight of the campaign.

To better understand how the optimal bid as described in Theorem 5.1 is obtained from the solution to a multi-agent non-cooperative game, consider the four main parts of the bid calculation (aka, players): an estimator of $\bar{v}_{\ell,i}$, a feedback controller estimating $\bar{\lambda}_0$, a feedback controller estimating $\bar{\lambda}_\ell$ for $\ell = 1, \ldots, \ell_{max}$, and an optimizer solving $b^{opt}_{\ell,i} = \arg\max_{0\leq b\leq b_{max}} \mathrm{E}\tilde{S}_{\ell,i}(b; b^u_i)$. Each player solves a specific sub-problem different from (7)-(10), without coordination with other players. However, if all players solve their own problem adequately, then the necessary conditions for optimality of the original problem are satisfied. In general, (7)-(10) is a non-convex problem; however, in many practical scenarios with only a modest number of constraints, any bidding strategy satisfying the necessary conditions is optimal or near-optimal. This result and the sufficiency conditions for optimal bidding is outside the scope of this paper.

Figure 1 provides a visual depiction of the interconnected

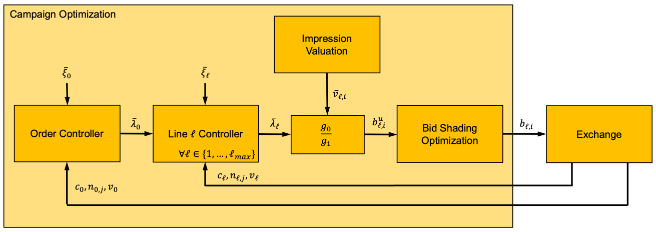

Fig. 1. The interconnected optimization system shown as a block diagram to illustrate how the subsystems are interconnected to compute a bid price $b_{\ell,i}$ for each impression opportunity. Signals $\bar{\xi}_0$, $\bar{\xi}_\ell$, $\bar{\lambda}_0$, $\bar{\lambda}_\ell$, and , $\bar{v}_{\ell,i}$ are vector-valued.

system involving the four types of players. Pay attention to the feedback loops indicating the algorithmic components that are naturally implemented as feedback controllers. The block diagram demonstrates how, in a modularized fashion, to compute a bid price $b_{\ell,i}$ for each impression opportunity. The practical value of the result is how it justifies implementing the bidding strategy across order and lines as non-cooperating subsystems.

A. Impression Valuation

This player of the game consumes historical impression data and corresponding context information, as well as, labels indicating the realized impression value. The information is used to train a prediction model for incoming new impression opportunities. The training typically takes place offline, e.g. once per day, using logistic regression, deep learning, or some other big data machine learning technique; whereas, computing the prediction of $\bar{v}_{\ell,i}$ for a specific impression opportunity is implemented via a real time scoring scheme based on context information available in the impression request before a bid is submitted [3], [4], [5].

B. Order ($\ell = 0$) Controller

This player ingests the order-level constraint vector $\bar{\xi}_0$, and order-level delivery feedback consisting of time-series observations of c_0 and $n_{0,j}$ for all j. The objective is to update $\bar{\lambda}_0$ so that it converges to its optimal value, where the corresponding Lagrangian conditions are satisfied. Note that feedback from individual impressions or lines is not needed (illustrating the non-cooperation). The controller is typically implemented as an adaptive discrete-time error feedback controller. Adaptation is needed since the effective plant gain (and potentially also the plant delay/dynamics) depends on environmental factors and varies across campaigns and over time. The system may be identified using e.g. as a standard recursive least squares estimator [10], [15]. A layman interpretation of $\bar{\lambda}_0$ is as a vector of break pedals. If a constraint is violated it typically means the corresponding element of $\bar{\lambda}_0$ needs to increase, and if the constraint is not binding but larger than zero, then it needs to decrease.

C. Line ℓ Controller

This player operates similar to order controller, but is based on line ℓ data (constraints $\bar{\xi}_\ell$ and feedback c_ℓ, $n_{\ell,j}$ for all j) to update $\bar{\lambda}_\ell$ and is subject to lower bounds $\bar{\lambda}_\ell \geq \bar{\lambda}_0$. That is, line ℓ controller seeks to find the value of $\bar{\lambda}_\ell$ for which the corresponding Lagrangian conditions are satisfied. Similar to order controller, $\bar{\lambda}_\ell$ can be interpreted as a set of break pedals, and Theorem 5.1 states that no line should use less "break force" than order controller.

The alternative would be to estimate all of $\bar{\lambda}_0$ and $\bar{\lambda}_\ell$ for $\ell = 1, \dots, \ell_{max}$ simultaneously, which is a much higher dimensional control problem. Another consequence of Theorem 5.1 reducing possible complex dynamic interaction between order and line controller relates to the determinants of updates of $\bar{\lambda}_\ell$. If a particular line constraint is not binding, then the corresponding element of $\bar{\lambda}_\ell$ inherits the corresponding element in $\bar{\lambda}_0$ no matter how noisy the line level feedback signal is. Similarly, if line controller must adjust an element of $\bar{\lambda}_\ell$ above its lower bound in order to avoid a constraint violation, then updates to this element is determined fully by line level feedback no matter how noisy the order level feedback signal is and even if the corresponding element of $\bar{\lambda}_0$ is somewhat volatile.

D. Bid Shading Optimization

This player (subsystem) consumes historical bids i', ℓ, the corresponding adjusted impression value $b^u_{\ell,i'}$ and context information. It trains a prediction model that is used for each incoming impression opportunity to directly compute $b^{opt}_{\ell,i} = \text{argmax}_b\, \tilde{s}_i(b^u_{\ell,i}, b)$ [13], or indirectly via $b^{opt}_{\ell,i} = \text{argmax}_b (b^u_{\ell,i} - b) F_{B^\star_i}(b)$ [11], [12], [14], depending on what information is available for model training. Note, $F_{B^\star_i}(b)$ is typically not known ahead of time, but must for the indirect method be predicted based on historical data.

VII. Example

The goal is to apply Theorem 5.1 on a toy example to illustrate how a basic hierarchical multi-constraint optimization problem can be solved using non-cooperating components. It is not the intent to offer a comprehensive assessment or to make the setup particularly realistic.

A. Problem Setup

Consider an ad campaign defined by an order consisting of three lines. The campaign objective is to maximize the expected campaign-aggregate total value (e.g. the total value of sold goods triggered by ad impressions awarded to the campaign). The order and the three lines are subject to daily spend constraints ξ_0, ξ_1, ξ_2, and ξ_3; and are managed by a discrete-time control system with a sampling time of 1/60 hours. The implementation is defined by equidistant time points indexed $t = 1, 2, \dots, t_{max}$; where t_{max} is the number of time intervals in the optimization window (one day); hence $t_{max} = 24 \cdot 60 = 1440$. Mathematically, the problem is

$$\underset{\{b_{\ell,i} | \forall i \in \Omega^\ell, \forall \ell \in \{1,2,3\}\}}{\text{maximize}} \quad \mathrm{E}V_0$$

subject to $\mathrm{E}C_0 \leq \xi_0$, $\mathrm{E}C_1 \leq \xi_1$, $\mathrm{E}C_2 \leq \xi_2$, $\mathrm{E}C_3 \leq \xi_3$. The campaign flight is 12 days long and is broken down into four intervals, each with a different set of spend constraints:

Interval number	Time interval (hr)	Daily order budget	Daily line 1 budget	Daily line 2 budget	Daily line 3 budget
1	[0,72)	\$70	\$70	\$10	\$5
2	[72, 144)	\$60	\$20	\$10	\$30
3	[144, 216)	\$100	\$10	\$20	\$10
4	[216, 288]	\$50	\$30	\$30	\$20

The sum of the line budgets exceed the order budget in intervals 1 and 4, equals the order budget in interval 2, and is less than the order budget in interval 3. Assume solutions to impression valuation [3], [4], [5], and bid shading optimization [11], [12], [13], [14] are implemented and operational (see Figure 1). It remains to implement and commission the order and line controllers (one per line). It is reasonably easy to prove that $\mathrm{E}V_\ell$ and $\mathrm{E}C_\ell$ by necessity are non-increasing functions of $\bar{\lambda}_\ell \equiv \lambda_\ell \in \mathbb{R}$; and that the optimal solution corresponds to the smallest values of $\lambda_\ell \geq \lambda_0 \geq 0$ for which no constraint is violated.

B. Control Problem and Methodology

Let $\Omega^\ell(t)$ and $c_\ell(t)$ denote the number of impression opportunities and the realized spend for line ℓ in time interval t, respectively. The corresponding value for the order are $\Omega^0(t) = \sum_{\ell=1}^{\ell_{max}} \Omega^\ell(t)$ and $c_0(t) = \sum_{\ell=1}^{\ell_{max}} c_\ell(t)$. Summing $\Omega^\ell(t)$ and $c_\ell(t)$ over t from 1 to t_{max} yield Ω^ℓ and c_ℓ, respectively. The values of $\Omega^\ell(t)$ and $c_\ell(t)$ are a priori unknown, and $c_\ell(t)$ depends on $\lambda_\ell(t)$, which is the estimate of λ_ℓ used to compute bids in time interval t.

The objective is to adjust $\lambda_0(t) \geq 0$ and $\lambda_\ell(t) \geq \lambda_0(t)$, for $\ell = 1, 2, 3$, toward the smallest constant values for which the constraints (8)-(10) are not violated. Define error signal $e_\ell(t)$, for $\ell = 0, 1, 2, 3$, where $e_\ell(t) = \xi_\ell / t_{max} - c_\ell(t)$. Constraints (8)-(10) are satisfied if $\sum_t e_\ell(t)$ for all ℓ are non-negative. Order and line level controllers achieve this by using error feedback control to update $\lambda_\ell(t)$.

The plant is defined by the map from $\lambda_\ell(t)$ to $c_\ell(t)$, for $\ell = 0, 1, 2, 3$; which in real applications is nonlinear, dynamic, time-varying, and stochastic [10]. Typically, the plant is approximately linear near each operating point.

C. Control Design

Inspired by the online dual decomposition method, consider pure integral (I) error feedback control to update control signals $\lambda_\ell(t)$ for each ℓ separately. Recall that $\lambda_0(t) \geq 0$ and $\lambda_\ell(t) \geq \lambda_0(t)$, for $\ell = 1, 2, 3$. Since $\lambda_\ell(t)$ depends on $\lambda_0(t)$ it is natural to update $\lambda_0(t)$ before updating $\lambda_\ell(t)$. Integrator wind-up protection is obtained automatically by using $\lambda_\ell(t)$, for $\ell = 0, 1, 2, 3$, as the states of the controller and updating them as

$$\lambda_0(t) = \max\left(\lambda_0(t-1) - k_0 e_0(t), 0\right), \quad (14)$$
$$\lambda_\ell(t) = \max\left(\lambda_\ell(t-1) - k_\ell e_\ell(t), \lambda_0(t)\right), \ \ell = 1, 2, 3, \quad (15)$$

where $k_0, k_\ell \in \mathbb{R}_{>0}$ are the controller I gains selected to yield a loop gain equal to 0.05 for each controller. Hence, the I gains depend on the plant gains, which may be estimated online. This estimation is outside the scope of this paper, but two approaches are discussed in [10].

D. Plant Model

The static plant of line ℓ is defined by the relationship between λ_ℓ and $\mathrm{E}C_\ell$. For the sole purpose of simulating a closed loop system, assume $\mathrm{E}C_\ell$, on a per day basis, can be described by $\mathrm{E}C_\ell = \gamma_\ell \Gamma(1/\lambda_\ell | \alpha_\ell, \beta_\ell)$, where γ_ℓ is the daily spend potential of the line, and $\Gamma(u|\alpha, \beta)$ denotes the CDF of a gamma distribution parameterized in terms of a shape parameter α and an inverse scale parameter β. This model is used of convenience for the simulation study to capture basic properties of the real plant (non-negative, monotonic, and bounded). The detailed plant model is irrelevant to the control system since the only thing of importance is the plant gain, which is estimated online (see above).

Parameters γ_ℓ, α_ℓ, and β_ℓ encode all relevant information related to $v_{\ell,i}$, $\Omega^\ell(t)$, and $F_{B_i^*}(b)$; which define impression value, impression supply, and competitive landscape. The line-specific parameters are provided in the following table:

Line ℓ	γ_ℓ	α_ℓ	β_ℓ
1	\$100	2	0.5
2	\$100	2	0.12
3	\$100	2	0.1

All three lines have the same overall spend potential (equal to \$100), but line 1 is able to spend more budget at larger values of λ_ℓ than lines 2 and 3 can. It can be shown that an ad budget that is spent when λ_ℓ is large yield impressions for which $v_{\ell,i}/c_{\ell,i}$ is large. Such impressions are referred to as high *return on investment* (ROI) impressions. The larger the ROI of awarded impressions, the more does each dollar in ad spend contribute towards the campaign objective. Although all three lines have access to both low and high ROI impressions, line 1 is able to spend more budget on high ROI impressions than what lines 2 and 3 are capable of.

The above relationships are known only to an oracle, and the control system must instead rely on observations $c_0(t)$ and $c_\ell(t)$ to update $\lambda_0(t)$ and $\lambda_\ell(t)$. These observations are assumed available at negligible system delay. The control system is implemented with controller I gains selected using an adaptive scheme to maintain a desired loop gain, stability, and acceptable performance. The details of the adaptation and gain selection are outside the scope of the paper. Moreover, the controller is initialized by $\lambda_0(0) = \lambda_1(0) = \lambda_2(0) = \lambda_3(0) = 0.1$.

E. Simulation Results

Figure 2 shows the closed loop result of a noise-free simulation of the above set-up. The left subplots show the order- and line-level control signals λ_0, λ_1, λ_2, and λ_3 versus time, and the right subplots show the order- and line-level spend and budget versus time. Budget and spend numbers in the plots are shown per sampling interval (one minute).

Note, how no budget constraint at steady-state is ever violated; i.e. $c_0(t) \leq \xi_0/t_{max}$ and $c_\ell(t) \leq \xi_\ell/t_{max}$, for $\ell = 1, 2, 3$. Next, consider each time interval separately.

In interval 1 ($0 \leq t < 72$), $\xi_1 + \xi_2 + \xi_3 \geq \xi_0$, which means at least one line cannot deliver its budget in full. In order for the Lagrangian condition $(\lambda_\ell - \lambda_0)(\mathrm{E}C_\ell - \xi_\ell) = 0$ to hold for that (one or more) ad line, it is necessary that $\lambda_\ell = \lambda_0$. We confirm from the figure that Line 1 exhibits the property of not spending the budget in full and that $\lambda_1 = \lambda_0$. Moreover, since the order-level spend potential is larger than ξ_o, the order level control signals $\lambda_0 \approx 0.27$ settles on a strictly positive number to prevent spending too much on the order-level. Line 2 and 3 spend their budgets in full using $\lambda_2 \approx 0.94$ and $\lambda_3 \approx 1.41$ strictly larger than λ_0.

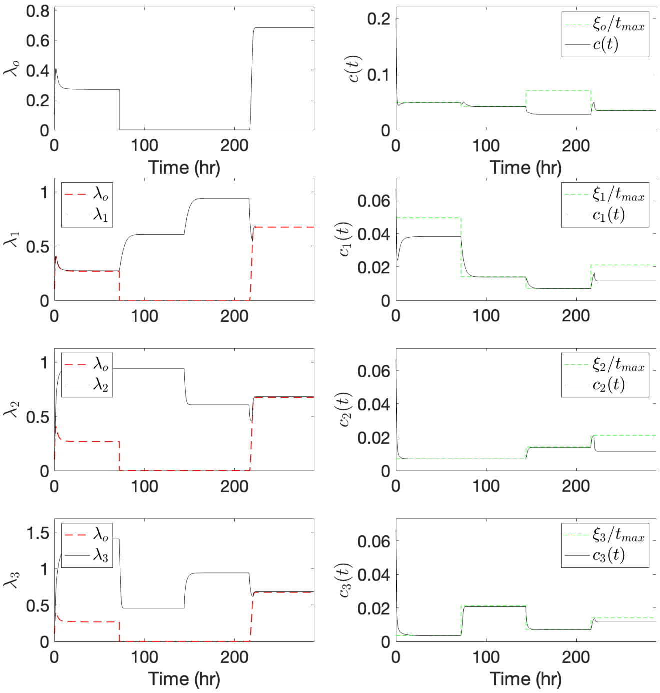

Fig. 2. Closed loop simulation result. Left: Control signals λ_0, λ_1, λ_2, and λ_3. Right: Spend signals $c_0(t)$, $c_1(t)$, $c_2(t)$, and $c_3(t)$; and budgets. Top row: Order result. Rows 2-4: Line 1-3 results.

In interval 2 ($72 \leq t < 144$), $\xi_1 + \xi_2 + \xi_3 = \xi_0$, which means cross-line budget optimization is not possible. Hence, the line-level controllers ensure the order level spend constraint is never violated. Each line controller manages its own objective with no concern about the other lines. Since all three lines have spend potentials that exceed their budgets, there exist values $\lambda_1 \approx 0.60$, $\lambda_2 \approx 0.94$, and $\lambda_3 \approx 0.46$ for which the lines spend their budgets in full. But since the order budget equals the sum of line budgets, this implies that the order budget is also spent in full.

In interval 3 ($144 \leq t < 216$), the constraints satisfy $\xi_1 + \xi_2 + \xi_3 \leq \xi_0$, which again means cross-line budget optimization is impossible. The line controllers ensure the order constraint is never violated and λ_0 therefore stays at, or converges to, zero. Each line controller manages its own objective and since their spend potentials exceed their budgets, there exist values $\lambda_1 \approx 0.94$, $\lambda_2 \approx 0.61$, and $\lambda_3 \approx 0.94$ for which the lines spend their budgets in full. However, the order budget exceed the sum of line budgets, consequently, the order leaves budget on the table.

In interval 4 ($216 \leq t \leq 288$), where $\xi_1 + \xi_2 + \xi_3 \geq \xi_0$, it turns out no line is able to deliver their budget in full, and for that reason $\lambda_1 = \lambda_2 = \lambda_3 = \lambda_0$. On the other hand, since the spend potential of the order exceeds the order budget $\lambda_0 \approx 0.68$ settles on a strictly positive value.

VIII. Conclusions and Future Work

Advertisers often define an ad campaign by an order consisting of multiple lines with campaign delivery constraints imposed hierarchically across order and lines. We have derived necessary conditions for the optimal bidding strategy that solves a large family of constrained optimization problems. The results show how the bidding strategy can be implemented as non-cooperating feedback control-based agents on the order and on individual lines. This reduces the number of control signals each individual feedback controller must compute and greatly simplifies the control design.

Real ad campaign plants are uncertain and approximately 24 hour periodic, and they are subject to significant non-additive stochastic noise [10]. The immediate next steps is therefore to develop algorithms for multi-input multi-output system identification. Future work also includes designing multi-input multi-output feedback controllers that can be used on the order and line level of a campaign.

References

[1] 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.

[2] 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.

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

[4] 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.

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

[6] 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.

[7] 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.

[8] 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.

[9] 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.

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

[11] 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.

[12] 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.

[13] 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.

[14] 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 *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'21), August 14-18, 2021, Singapore*, 2021.

[15] Niklas Karlsson. Scalable multi-objective optimization in programmatic advertising via feedback control. In *Proceedings of the 2021 60th IEEE Conference on Decision and Control (CDC), Austin, TX, USA*, pages 1363–1370, Dec 2021.

[16] Niklas Karlsson. Feedback control-based multiobjective optimization in programmatic advertising involving a cost per bid constraint. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 505–510, 2022.

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

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