Control of Periodic Systems in Online Advertising

Niklas Karlsson[1]

***Abstract*— We propose a control system to regulate a periodic plant subject to significant load disturbances and measurement noise and a highly uncertain plant gain, but with negligible dynamics between control input and output. This is a common control problem in online advertising. The controller is periodic involving a periodic feed-forward adjustment of the reference setpoint and a periodic *proportional-integral* (PI) feedback controller. We derive the solution of the closed-loop system and prove that it is globally asymptotically stable. The proposed controller is compared with a control system based on the above feed-forward component and a standard linear and time-invariant PI controller.**

I. Introduction

Programmatic advertising is a big business and is technically defined as the use of machines and algorithms to purchase ad display space. It is the cornerstone of the business model for companies such as Google, Facebook, and Oath. The strategy is to use data about users to serve relevant ads that drive revenue to advertisers and create a positive experience to Internet users. EMarketer's estimate on worldwide media spend in 2017 is nearly $600 billion [1], which is a number that speaks for itself. A *Demand Side Platform* (DSP) is a particular business model for programmatic advertising. It is the middleman between an advertiser and one or more open exchanges trading so called ad impressions. An *impression* is an opportunity of showing an ad creative, e.g., banner ad, text ad, or pre-roll video commercial to Internet users. The goal of a DSP is to manage an advertiser's advertisement budget optimally.

An early pre-DSP publication on feedback control applied to online advertising is available in [2], wherein several important challenges are outlined but detailed solutions are omitted. A more comprehensive and up-to-date overview of the control problem is available in [3]. There the author proposes a bid randomization technique [4] to overcome some of the challenges. One of the challenges is the need of estimating extremely small event rates [5], and another is of controlling event rates in discontinuous plants [6]. The fact that the plant is unknown, dynamic, nonlinear, and in general discontinuous is a characteristic property of online advertising processes and is a fundamental challenge in the development of feedback control solutions. Different approaches at indirectly estimating and controlling the plant are proposed in [7], [8]. The former of these papers makes use of bid randomization, and both are based exclusively on local feedback information to estimate the plant gain. Adaptive estimation and control is always challenging since it requires persistent excitation [9], but is particularly difficult when the unknown plant is nonlinear and potentially discontinuous. A systematic methodology for off-line modelling of advertising plants is proposed in [10]. The methodology is conveniently used to simulate realistic plants in a test bed for control algorithms. Statistical inference of the plant gain based on an uncertain discontinuous plant response curve combined with bid randomization is presented in [11]. Several theorems are proven that translate a nominal plant model, model uncertainty, and a known bid randomization level into statistical properties of the plant gain.

In this paper we go to the heart of the optimization problem and propose a periodic controller to regulate the spend rate of an advertising campaign subject to a budget constraint. We assume a periodic and locally linear plant, and prove global asymptotic stability of the closed-loop system.

The paper is organized as follows. A few well-known key results about time-varying systems are reviewed in Section II. Thereafter, in Section III we define the ad optimization problem and explain how it is turned into a control problem. Next, in Section IV the plant model is defined mathematically. The candidate controller structure is introduced in Section V, and is analyzed in the context of the closed-loop system in Section VI. Simulation results are provided in Section VII. Finally, some concluding remarks and planned future work are given in Section VIII.

II. Linear Time-varying Systems

A linear homogeneous time-varying system is defined by

$$\dot{x} = A(t)x, \tag{1}$$

where $x \in \mathbb{R}^n$, and $A(t)$ is continuous in time. The solution of (1) is given by $x(t) = \Phi(t,t_0)x(t_0)$, where $\Phi(t,t_0)$ is the state transition matrix, which is unique (see e.g. [12]) and may be computed using the infinite *Peano-Baker series*:

$$\begin{aligned}\Phi(t,t_0) &= I + \int_{t_0}^{t} A(\tau_1)d\tau_1 + \int_{t_0}^{t} A(\tau_1)\int_{t_0}^{\tau_1} A(\tau_2)d\tau_2 d\tau_1 \\ &\quad \cdots \\ &\quad + \int_{t_0}^{t} A(\tau_1)\int_{t_0}^{\tau_1} A(\tau_2)\cdots\int_{t_0}^{\tau_{m-1}} A(\tau_m)d\tau_m d\tau_{m-1}\cdots d\tau_1 \\ &\quad + \cdots\end{aligned}$$

If $A(t)$ and $\int_{t_0}^{t} A(\tau)d\tau$ commute; i.e., if $A(t)\int_{t_0}^{t} A(\tau)d\tau = \int_{t_0}^{t} A(\tau)d\tau A(t)$, then

$$\begin{aligned}\Phi(t,t_0) &= \exp\int_{t_0}^{t} A(\tau)d\tau \\ &:= I + \sum_{k=1}^{\infty}\frac{1}{k!}\left[\int_{t_0}^{t} A(\tau)d\tau\right]^k.\end{aligned}$$

[1] N. Karlsson is Chief Scientist and Vice President of Demand Platforms R&D at Oath Inc, 741 1st Ave., Sunnyvale, CA 94089, USA niklas.karlsson at oath.com

Note, if $A(t)$ is a scalar function, or a constant matrix, then $A(t)$ and $\int_{t_0}^{t} A(\tau)d\tau$ always commute.

Consider the linear nonhomogeneous time-varying system

$$\dot{x} \quad = \quad A(t)x + g(t), \tag{2}$$

subject to initial condition $x(t_0) = x_0$, where $x \in \mathbb{R}^n$; and where $A(t)$ and $g(t)$ are continuous in time and of the obvious dimensions.

Theorem 2.1: Let $\Phi(t, t_0)$ denote the state transition matrix for (1) for all t. Then the unique solution $\phi(t, t_0, x_0)$ of (2) satisfying $\phi(t_0, t_0, x_0) = x_0$ is given by

$$\phi(t, t_0, x_0) \quad = \quad \Phi(t, t_0)x_0 + \int_{t_0}^{t} \Phi(t, \tau)g(\tau)d\tau.$$

Proof: See [12]. ■

We have glossed over some details concerning existence, uniqueness, continuation of solution, etc., but this is standard material readily available in the literature [12], [13], [14]. In particular, the explicit solution either is still valid or can be modified in an obvious way in the event that matrix $A(t)$ and vector $g(t)$ are only piecewise continuous rather than continuous.

III. Problem Formulation

The impression allocation for an impression i is governed by a sealed second price auction, where b_i is the bid price submitted by the DSP on behalf on an advertiser and used in the auction. Denote the highest competing bid price $b_i^\star$, and assume the advertiser is interested in some *event* with a monetary value of v_E. For example, an impression, a click, a conversion, or a combination of some of the above. An event is conditioned on an impression first taking place and the probability of an event to occur is then p_i. Note, if the event of interest to the advertiser is an impression, then $p_i \equiv 1$.

Define the total value v of an ad campaign as the expected number of events generated from the ad spend and c as the cost spent for the impressions. An impression is awarded to the advertiser if $b_i \geq b_i^\star$, and the advertiser is then charged $b_i^\star$ for the impression. It follows that

$$v \quad = \quad \sum_i p_i v_E \mathbb{I}_{\{b_i \geq b_i^\star\}} \tag{3}$$

$$c \quad = \quad \sum_i b_i^\star \mathbb{I}_{\{b_i \geq b_i^\star\}} \tag{4}$$

where $\mathbb{I}_X$ is the indicator function satisfying $\mathbb{I}_X = 1$, if $X =$ true, and $\mathbb{I}_X = 0$, if $X =$ false. It was shown in [3] that the maximization of v subject to a budget constraint $c \leq u_{cost}^{ref}$ is achieved using a bidding strategy $b_i = up_i$, where u is selected as the largest value for which the budget constraint is not violated. This result is illustrated graphically in Figure 1. Each impression we bid for is associated with an event rate p_i and a cost $b_i^\star$, hence can be mapped to a coordinate in the event rate versus cost plot. Impressions in the upper left corner with high event rate and low cost correspond to the highest *return on investment* (ROI) impressions since they translate to the largest number of *events* per ad dollar. Impressions along a straight line going through the origin

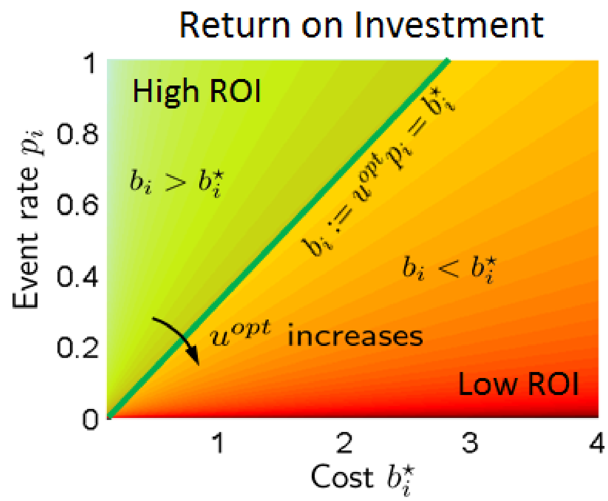

Fig. 1. The plot shows the relative ROI of different impression opportunities, where each impression is associated with an event rate p_i and a cost $b_i^\star$. High ROI impressions are located in the upper left corner, and low ROI impressions in the bottom right.

all have the same ROI, and impressions in the lower right corner with low event rate and high cost correspond to the lowest ROI impressions. Finally, impressions in the green shaded region all have higher ROI than impressions in the orange shaded region. By bidding along the straight line, we win the impressions in the green region and the optimization turns into a control problem where the objective is to choose the angle of the line, or u, for which the budget is spent but the constraint is not violated.

The delay from when we compute and use u and b_i to when we know whether we won the auction and how much we are charged for an awarded impression is short. The delay is typically in the order of minutes, but there are still several challenges. First, there is a dramatic deterministic seasonality in the number of available impressions. The seasonality is primarily the result of the periodicity in Internet traffic. The most dramatic seasonality is the time-of-day pattern. At the daily high there are about ten times more people online compared to near the daily low. Furthermore, Internet traffic contains a significant level of stochastic noise. Moreover, the available impressions are not distributed uniformly in the event rate versus cost plot, but in some a priori unknown clusters resulting in a discontinuous, staircase-like relationship between u and c (or v). The latter feature is illustrated in Figure 2, where the awarded number of impressions as a function of u is shown for two real and representative ad campaigns. Note how Campaign A is well-behaved in the sense that it, from a feedback control perspective, represents an almost continuous plant. In contrast, Campaign B represents a challenging discontinuous plant. We shall not discuss the matter in detail in this paper, but in situations where the discontinuity cannot be neglected, the effective relationship can be turned continuous by leveraging Heisenberg bidding [3], which is a bid randomization technique with which a computed nominal bid b_i is perturbed slightly before

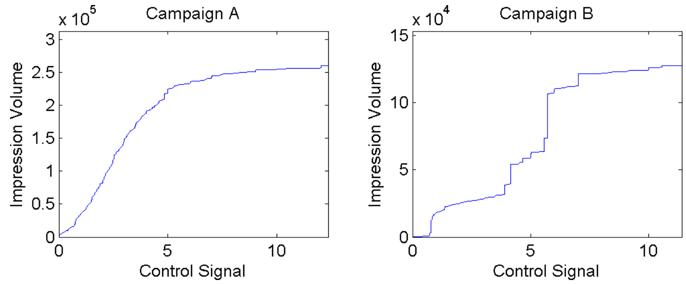

Fig. 2. Two representative examples of awarded number of impressions versus control signal u. Campaign A is a well-behaved campaign with an approximately smooth relationship while Campaign B is a challenging campaign with pronounced large steps in the relationship.

submitted to the auction exchange. Depending on the value of the so called uncertainty signal (or uncertainty bid), the input-output relationship of the plant can be made arbitrarily smooth. This is illustrated in Figure 3, which shows the result

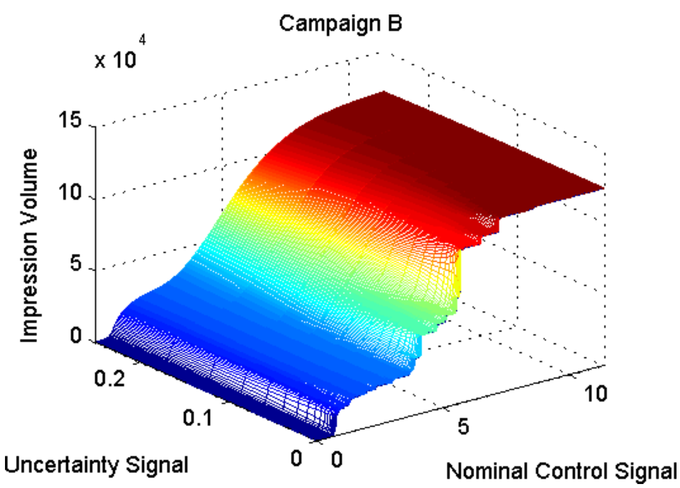

Fig. 3. The impact of using bid uncertainty on Campaign B in Figure 2. Note that the nominal control signal versus impression volume relationship is discontinuous only when the uncertainty signal equals zero.

of adding the dimension of uncertainty signal to Campaign B in the previous figure.

For the remainder of this paper we assume the relationship between u and y is continuous and that the slope at each operating point is highly uncertain. The slope is referred to as plant gain and is denoted K_p.

IV. Plant Model

Consider a plant which locally around the operating point can be described by a linear time-periodic model having insignificant dynamics but dramatic seasonality. The seasonality is in form of the T-periodic plant gain $K_p(1+h(t)) > 0$, where $h(t) \in C$, $h(t) > -1$, and $\int_t^{t+T} h(\tau)d\tau = 0$. The plant maps a bid adjustment control signal u to an ad spend rate y. Assume the system is subject to a load disturbance $v_\ell(t)$ and measurement noise $v_m(t)$ entering the plant as indicated in Figure 4. Measurement noise is here interpreted as the random deviation of the spend rate away from its expected value and is caused by the stochastic

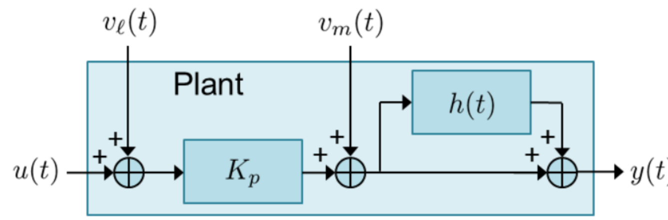

Fig. 4. Block diagram of the plant model.

behavior of Internet traffic. Note, by definition $\mathrm{E}\, v_m = 0$. Load disturbance is the reflection of a dynamic competitive landscape among advertisers – competing advertisers adjust their bids up and down over time.

The input-output relationships of the plant are mathematically described by

$$y = \big(1+h(t)\big)\Big(K_p\big(u+v_\ell\big)+v_m\Big), \tag{5}$$

We may model the $u \mapsto y$ relationship as non-dynamic by ensuring the dynamics of the control system is dominant. Plant gain K_p may evolve dynamically, but this dynamic is in general slow and is disregarded in this paper.

V. Controller

The main challenges of the control problem is that plant gain K_p is highly uncertain and difficult to estimate, and the measurement noise is significant. It is therefore necessary to design a relatively conservative controller to avoid dramatic fluctuations in the control signal. The seasonal component of the plant, $h(t)$, on the other hand, can be reasonably well estimated based on historical data. This fact is taken advantage of in the forthcoming proposed control design.

Assume the periodic component of the plant $h(t)$ is known, and the set-point signal $\bar{u}_c(t)$ is the specified average spending rate over time period T (e.g. one day). The objective is to design an error feedback controller that maximizes the ROI for the advertiser. This is achieved as explained in Section III by finding the single value of bid adjustment u that results in a spend rate that on average over a cycle, say, day equals $\bar{u}_c(t)$. With an optimal value of u and a fixed competitive landscape, the ad budget is spent in accordance with the plant seasonality $h(t)$ with more budget spent during hours of the day when more people are online and less in hours when few people are online.

To distribute the budget across the day we use a *linear time-periodic* (LTP) feed-forward adjustment of the set-point signal and define the feedback error by

$$e = \big(1+h(t)\big)\bar{u}_c - y. \tag{6}$$

Consider the following first order linear time-varying feedback controller

$$\dot{x} = a(t)x + b(t)e \tag{7}$$

$$u = c(t)x + d(t)e \tag{8}$$

The controller structure is illustrated in Figure 5, where the feedback controller is depicted by $f(t,e)$. We defer

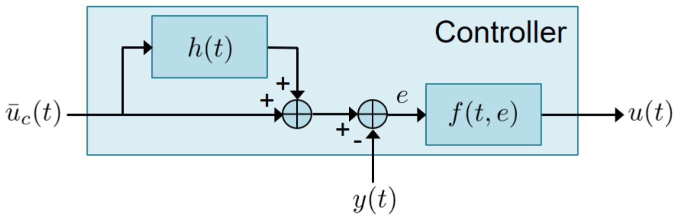

Fig. 5. Block diagram of the controller.

to the next section for the specific choice of the time-varying functions $a(t), b(t), c(t)$ and $d(t)$, since the choice goes hand-in-hand with a careful analysis of the closed-loop dynamics.

VI. Closed-Loop System

To determine the closed-loop dynamics, let us first derive the mapping from the controller state x and exogenous input signals $\bar{u}_c, v_\ell, v_m$ to feedback error e; i.e., $(x, \bar{u}_c, v_\ell, v_m) \mapsto e$. Combine expressions (5) and (6) to obtain

$$e = (1+h(t))\left(\bar{u}_c - K_p(u+v_\ell) - v_m\right).$$

Replacing u in this equation by (8) leads to

$$e = (1+h(t))\left(\bar{u}_c - K_p(c(t)x + d(t)e + v_\ell) - v_m\right).$$

Solve this equation for e and assume $1 + K_p(1+h(t))\,d(t) \neq 0$, we have

$$\begin{aligned} e &= \frac{-K_p(1+h(t))c(t)}{1+K_p(1+h(t))\,d(t)}x \\ &\quad + \frac{(1+h(t))\left(\bar{u}_c - K_p v_\ell - v_m\right)}{1+K_p(1+h(t))\,d(t)}. \end{aligned} \tag{9}$$

The closed-loop dynamics of state x is obtained by combining (7) and (9) to be

$$\begin{aligned} \dot{x} &= \frac{a(t) + K_p(1+h(t))\left(a(t)d(t) - b(t)c(t)\right)}{1+K_p(1+h(t))\,d(t)}x \\ &\quad + \frac{b(t)(1+h(t))(\bar{u}_c - K_p v_\ell - v_m)}{1+K_p(1+h(t))\,d(t)} \end{aligned} \tag{10}$$

We are interested in knowing if in expected sense there exists a stable equilibrium of x when $\bar{u}_c$ and v_ℓ are constant. Since v_m enters the dynamics in a linear and additive fashion and $\mathrm{E}\, v_m = 0$, we may set $v_m = 0$ to explore the expected dynamics. A necessary condition for an equilibrium to exist is that there is a value of $\mathrm{E}\, x$ at which $\mathrm{E}\, \dot{x} = 0$. Denote this value $x^\star$. It follows that

$$x^\star = -\frac{b(t)(1+h(t))(\bar{u}_c - K_p v_\ell)}{a(t) + K_p(1+h(t))\left(a(t)d(t) - b(t)c(t)\right)},$$

where we have assumed the denominator is not equal to zero. In order for $x^\star$ to be an equilibrium it must be independent of t. This is achievable by a proper choice of the time-varying parameters of the feedback controller. A particularly good choice, as will be shown, is $a(t) = 0$, $b(t) = b = \text{const} > 0$, $c(t) = c = \text{const} > 0$, and $d(t) = c/(1+h(t))$. This represents a feedback controller described by

$$\dot{x} = be \tag{11}$$

$$u = cx + \frac{c}{1+h(t)}e \tag{12}$$

Note, (11) and (12) defines a time-periodic PI controller with a constant integral gain bc and a time-periodic proportional gain $c/(1+h(t))$. The closed-loop dynamics then becomes

$$\begin{aligned} \dot{x} &= -\frac{K_p(1+h(t))bc}{1+K_pc}x \\ &\quad + \frac{b(1+h(t))(\bar{u}_c - K_p v_\ell - v_m)}{1+K_pc} \end{aligned} \tag{13}$$

The selected feedback controller leads to a closed-loop system where $x^\star$ is the unique equilibrium in expected sense and given by

$$x^\star = \frac{1}{K_pc}\left(\bar{u}_c - K_p v_\ell\right).$$

To determine whether $x = x^\star$ is stable in expected sense, compute the expectation of (13):

$$\mathrm{E}\,\dot{x} = -\frac{K_pbc(1+h(t))}{1+K_pc}\mathrm{E}\,x + \frac{b(1+h(t))(\bar{u}_c - K_p v_\ell)}{1+K_pc} \tag{14}$$

Define $\tilde{x} = x - x^\star$ as the deviation of x from the equilibrium and derive its dynamics in expected sense.

$$\begin{aligned} \mathrm{E}\,\dot{\tilde{x}} &= -\frac{K_pbc(1+h(t))}{1+K_pc}\left(\mathrm{E}\,\tilde{x} + x^\star\right) \\ &\quad + \frac{b(1+h(t))(\bar{u}_c - K_p v_\ell)}{1+K_pc} \\ &= -\frac{K_pbc(1+h(t))}{1+K_pc}\mathrm{E}\,\tilde{x} \end{aligned} \tag{15}$$

By inspection we observe that $\mathrm{E}\,\tilde{x} > 0 \Rightarrow \mathrm{E}\,\dot{\tilde{x}} < 0$ and $\mathrm{E}\,\tilde{x} < 0 \Rightarrow \mathrm{E}\,\dot{\tilde{x}} > 0$, which implies that $\mathrm{E}\,\tilde{x} \to 0$ as $t \to \infty$. In other words, $\mathrm{E}\,x \to x^\star$, and hence $x^\star = (\bar{u}_c - K_p v_\ell)/(K_pc)$ is a globally stable equilibrium of $\mathrm{E}\,x$ when $\bar{u}_c$ and v_ℓ are constant.

The closed-loop feedback error (9) with the selected controller parameters satisfies

$$\begin{aligned} e &= -\frac{K_p(1+h(t))c}{1+K_pc}x \\ &\quad + \frac{(1+h(t))(\bar{u}_c - K_p v_\ell - v_m)}{1+K_pc}. \end{aligned} \tag{16}$$

Meanwhile, control signal u defined by (12) with e replaced by (16) satisfies

$$u = \frac{c}{1+K_pc}x + \frac{c(\bar{u}_c - K_p v_\ell - v_m)}{(1+K_pc)} \tag{17}$$

Furthermore, the output of the plant y is defined by (5) and with u given by (17) becomes

$$y = \frac{1+h(t)}{1+K_pc}\left(K_pcx + K_pc\bar{u}_c + K_p v_\ell - v_m\right).$$

As $t \to \infty$ we have shown that $\mathrm{E}\, x \to (\bar{u}_c - K_p v_\ell)/(K_p c)$. Using this in the expressions for e, u, and y, together with some simplifications, we obtain

$$\begin{aligned} \mathrm{E}\, e &\to 0 \\ \mathrm{E}\, u &\to \frac{\bar{u}_c}{K_p} - v_\ell \\ \mathrm{E}\, y &\to (1+h(t))\bar{u}_c. \end{aligned}$$

Hence, whenever $\bar{u}_c$ and v_ℓ are constant it follows that the expected values of e and u converge to stable equilibria and the expected value of y converges to a stable limit cycle. Note also that the average expected spend rate y_{ave} in any time interval $[t, t+T)$ is given by

$$\begin{aligned} y_{ave} &= \frac{1}{T}\int_t^{t+T} y(\tau) d\tau \\ &\to \frac{1}{T}\int_t^{t+T} (1+h(\tau))\bar{u}_c d\tau \\ &= \bar{u}_c, \end{aligned}$$

which is in accordance with the control requirement.

Let us now determine the trajectory of $\mathrm{E}\, x$ for arbitrary trajectories of $\bar{u}_c$ and v_ℓ. A direct application of Theorem 2.1 on (14) implies that the solution $\mathrm{E}\, x(t)$ for $x(t_0) = x_0$ is given by

$$\phi(t, t_0, x_0) = \Phi(t, t_0)x_0 + \int_{t_0}^{t} \Phi(t, \tau) g(\tau) d\tau, \quad (18)$$

where

$$\begin{aligned} \Phi(t, t_0) &= \exp\left(-\int_{t_0}^{t} \frac{K_p bc(1+h(\tau))}{1+K_p c} d\tau\right), \\ g(t) &= \frac{b(1+h(t))(\bar{u}_c(t) - K_p v_\ell(t))}{1+K_p c}. \end{aligned}$$

It is in general cumbersome to compute the integral function in (18), but in the special case of constant $\bar{u}_c$ and v_ℓ we may simplify the expression by applying the theorem on the trajectory of $\mathrm{E}\, \tilde{x} = \mathrm{E}\, x - x^\star$ defined by (15). The initial value is given by $\tilde{x}(t_0) = x_0 - x^\star$ and the solution $\mathrm{E}\, \tilde{x}(t)$ is obtained as $\mathrm{E}\, \tilde{x}(t) = \Phi(t, t_0)\tilde{x}(t_0)$. Consequently

$$\begin{aligned} \mathrm{E}\, x(t) &= \Phi(t, t_0)\tilde{x}(t_0) + x^\star \\ &= \Phi(t, t_0)\left(x_0 - x^\star\right) + x^\star \\ &= \Phi(t, t_0)x_0 + \left(1 - \Phi(t, t_0)\right) x^\star, \end{aligned}$$

where $\Phi(t, t_0)$ can be rewritten as

$$\Phi(t, t_0) = \exp\left(-\frac{K_p bc(t - t_0)}{1+K_p c} - \frac{K_p bc}{1+K_p c}\int_{t_0}^{t} h(\tau) d\tau\right).$$

We have computed the globally stable equilibrium of the closed-loop system, and devised a formula to compute the trajectory of the system towards the equilibrium.

To conclude this section, we suggest parameterizing the controller using $b = 1/T_i$ and $c = K_c$; i.e.,

$$\dot{x} = \frac{1}{T_i} e \quad (19)$$

$$u = K_c x + \frac{K_c}{1+h(t)} e, \quad (20)$$

which is an LTP PI controller in standard form [15] having a proportional term with a periodic gain.

VII. Simulation Results

Consider a plant defined by (5) with $K_p = 1$ and $h(t) = 0.8\sin(2\pi t/24 + 1.5\pi)$. We shall compare the behavior of the proposed LTP PI controller defined by (see also (6), (19), and (20))

$$\begin{aligned} e &= \big(1+h(t)\big)\bar{u}_c - y \\ \dot{x} &= \frac{1}{T_i} e \\ u &= K_c x + \frac{K_c}{1+h(t)} e, \end{aligned}$$

with a conventional PI feedback controller, which is *linear time-invariant* (LTI):

$$e = \big(1+h(t)\big)\bar{u}_c - y \quad (21)$$

$$\dot{x} = \frac{1}{T_i} e \quad (22)$$

$$u = K_c x + K_c e. \quad (23)$$

The comparison is done for the single set of controller parameters $K_c = 0.5$ and $T_i = 15$. Suppose the set-point signal is given by

$$\bar{u}_c(t) = \begin{cases} 1, & \text{if } t < 150 \\ 0.75, & \text{otherwise} \end{cases}$$

and assume the load disturbance is

$$v_\ell(t) = \begin{cases} 0, & \text{if } t < 250 \\ -0.5, & \text{if } 250 \le t < 350 \\ 0.5, & \text{otherwise} \end{cases}$$

We evaluate only the behavior in expected sense, hence set $v_m = 0$; but recognize that a comprehensive assessment of the control system and a tuning of controller parameters requires us to consider also measurement noise input, unmodelled dynamics, and model uncertainties, which is outside the scope of this paper. Indeed, in practice the seasonality model $\hat{h}(t)$ used in the controller will never be identical to the true plant seasonality $h(t)$.

The closed-loop behavior for each of the two feedback controllers in the time range $0 \le t \le 450$ is shown in Figure 6. Note how the steady-state behavior is the same for both control systems while the benchmark conventional LTI PI controller exhibits a significant oscillatory transient pattern visible primarily in the control signal u, but to some extent also in other signals.

The oscillatory control signal compromises the ROI of the ad campaign since unnecessarily large values of u at any time means lower quality impressions are purchased. However, a potentially more dramatic impact of oscillations is in light of the true plant being discontinuous and more accurately represented by a staircase function [3]. It easily leads to a highly volatile spend rate y and may trigger limit cycles in the closed-loop dynamics, or potentially even chaotic behavior.

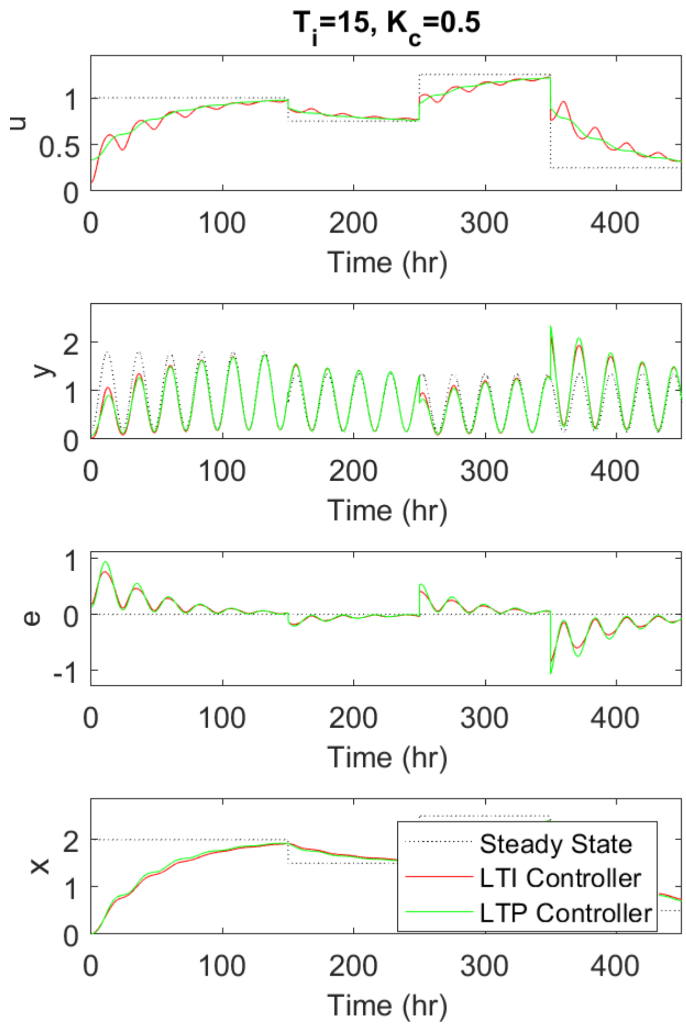

Fig. 6. Simulated closed-loop result for the proposed periodic PI controller (green-LTP) and a benchmark conventional PI controller (red-LTI). From top to bottom: control signal u, spend rate y, feedback error e, and state x.

With the caveat that the simulation results are limited, the proposed periodic controller appears to be superior over the conventional PI controller. Note, the periodic controller does not introduce additional dynamics. On the contrary, by introducing a time-varying controller parameter, the closed-loop dynamics is simpler.

VIII. Conclusions and Future Work

We have proposed a periodic controller for a plant with significant periodicity, load disturbances, and measurement noise, but with negligible dynamics. The closed-loop system is proven to be globally asymptotically stable, and the closed form solution of the state (and other signals) is derived. It is shown in simulations how the proposed controller outperforms a corresponding standard non-periodic feedback controller.

We recognize the limitations of the assessment completed in this paper. Important future work involves an assessment of the closed-loop system behavior in response to realistic load disturbances and measurement noise, and an evaluation of the system impact from unmodelled dynamics and modelling uncertainties. Note, in this paper we assumed the controller has a perfect knowledge of the plant seasonality $h(t)$. This may not be possible in practice, hence must be investigated further as part of future work. A comprehensive experimental evaluation is also an important next step.

References

[1] eMarketer. Worldwide ad spending: emarketer's updated estimates and forecast for 2016-2021.

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

[3] Niklas Karlsson. Control problems in online advertising and benefits of randomized bidding strategies. *European Journal of Control*, 30:31–49, July 2016.

[4] Niklas Karlsson. *Systems and Methods for Controlling Bidding for Online Advertising Campaigns, United States Patent Application 20100262497*. USPTO, 2009.

[5] Niklas Karlsson. Adaptive estimation of small event rates. *Proceedings of the 55th IEEE Conference on Decision and Control, Las Vegas, NV, USA, December 12 - 14*, pages 3732–3737, 2016.

[6] Niklas Karlsson and Qian Sang. Event rate control in online advertising. *Proceedings of the 2017 American Control Conference, Seattle, WA, USA*, pages 4215 – 4220, May 24-26 2017.

[7] Niklas Karlsson. Adaptive control using Heisenberg bidding. *Proceedings of the 2014 American Control Conference, Portland, USA, June 4-6, 2014*, pages 1304–1309, 2014.

[8] Jiaxing Guo and Niklas Karlsson. Model reference adaptive control of advertising systems. *Proceedings of the 2017 American Control Conference, Seattle, WA, USA*, pages 5482 – 5487, May 24-26 2017.

[9] Karl Johan Åström and Björn Wittenmark. *Adaptive Control*. Prentice Hall, second edition, 1994.

[10] Vahid Mardanlou, Niklas Karlsson, and Jiaxing Guo. Statistical plant modeling and simulation in online advertising. *Proceedings of the 2017 American Control Conference, Seattle, WA, USA*, May 24-26 2017.

[11] Niklas Karlsson. Plant gain estimation in online advertising processes. *Proceedings of the 2017 IEEE 56th Conference on Decision and Control, Melbourne, Australia*, pages 2182 – 2187, Dec 12-15 2017.

[12] Panos J. Antsaklis and Anthony N. Michel. *Linear Systems*. Birkhaüser, 2006.

[13] Kostas S. Tsakalis and Petros A. Ioannou. *Linear Time-varying Systems*. Prentice Hall, 1993.

[14] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, third edition, 2002.

[15] Karl Johan Åström and Tore Hägglund. *Advanced PID Control*. ISA – The Instrumentation, Systems, and Automation Society, 2005.