# Adaptive Control Using Heisenberg Bidding

**Proceedings of the 2014 American Control Conference, Portland, OR, June 4 - 6, 2014**

Niklas Karlsson[1]

***Abstract*— We propose an actuator mechanism to enable and improve scalable control of bidding based networks. The mechanism is a bid randomization strategy referred to as Heisenberg bidding, which is used to modify the effective plant gain of a system involving a large number of similar objects distributed via an auction exchange. The bid randomization mechanism is used in a time-varying model-reference adaptive controller and applied to a problem of satisfying traffic commitments in the content network of a major web portal.**

## I. INTRODUCTION

The explosive growth of networks on Internet has made the demand for highly scalable algorithms for robust and adaptive optimization and control greater than ever before. While the networks may represent vastly different systems e.g. display advertising networks, content delivery networks, social networks, shopping networks (Amazon, Netflix, etc.); they all define allocation optimization problems that are very high-dimensional, dynamic, and subject to tight response time requirements on allocation decisions.

It is popular to use an auction exchange with decentralized bidding strategies to solve the above type of problems. However, while this potentially leads to a scalable solution, it introduces challenges of a different nature. Indeed, it leads to discontinuous input-output relationships.

The *Heisenberg-bidding based control* method [1] introduced in this paper is a decentralized algorithm utilizing bid randomization implemented in an auction exchange. Each bidding agent relies only on measurement data related to itself. Related work and practical challenges in the area of ad and content optimization are described in [2], [3], [4].

The paper is organized as follows. A high-level overview of the system is provided in Section II. Thereafter, in Section III Heisenberg bidding is introduced. The control problem is defined in Section IV and a model of the plant is presented in Section V. Section VI derives one possible control algorithm making use of Heisenberg bidding. Some simulation results are shown in Section VII and some concluding remarks given in Section VIII.

## II. OVERVIEW

One possible implementation of Heisenberg bidding based control is shown in Figure 1. The diagram indicates how the key components of the algorithm are connected, and how an add-on component to support exploration & exploitation (not covered in this paper) can be introduced [5]. Without loss of generality we assume the objects to allocate are "impressions", e.g. display ads, content promotions, or product recommendations. The impressions are distributed via an auction exchange (Market Clearing in the block diagram) in real-time when an Internet user loads a web page and a decision is made which ad, content promotion, or product recommendation to display. For each available impression the auction exchange receives bids from all interested bidders, and a market clearing takes place to select which bidder is awarded the impression. Only the market clearing component handles information from multiple (all) bidders. All other components appear once per bidder and are bidder centric with no information exchange with other bidders.

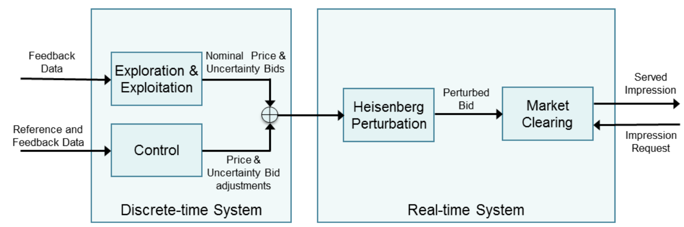


Fig. 1. Block diagram of Heisenberg bidding based Control.

The control component for each bidder is fed reference and feedback data, consisting of a desired rate of impressions (or clicks, revenue, etc.) per time unit and an observed number of impressions (or clicks, revenue, etc.) per sample interval. This information is used to calculate a bid price $u_p(t) \in \mathbb{R}$ and a bid uncertainty $u_u(t) \in \mathbb{R}$ adjustment. The price adjustment is the primary control lever, but is assisted by the uncertainty adjustment to modify the effective plant gain as needed. For example, it reduces the plant gain at discontinuities where the gain otherwise is $\infty$, and increases the plant gain in intervals where the gain otherwise is zero.

Again, exploration and exploitation is outside the scope of this paper, but would otherwise produce outputs that are combined with $u_p(t)$ and $u_u(t)$. In the sequel we assume exactly $u_p(t)$ and $u_u(t)$ are fed to Heisenberg Perturbation in which a perturbed, so called, *final bid price* is generated by random draw. The final bid price is used in a standard market clearing in which the bidder submitting the highest bid price is awarded the impression.

The random draw and the market clearing takes place once per auction, while the calculation of $u_p(t)$ and $u_u(t)$ typically happen only in discrete time.

## III. HEISENBERG BIDDING

The idea of Heisenberg bidding was first published in [6]. It introduces into auctions an artificial analogue of the quantum mechanical concepts of uncertainty principle and

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

tunneling. Indeed, before clearing the market Heisenberg bidding randomly perturbs each submitted *nominal bid price* $u_p$ based on a submitted *bid uncertainty* $u_u$ to generate a *final bid price* $B$ used in the market clearing. Heisenberg bidding can be implemented with other probability distributions, but in this paper it is defined by $B \sim \text{Gamma}(\alpha, \beta)$ where $\alpha > 0$ and $\beta > 0$ are the shape and inverse scale parameters of the Gamma distribution, and where we let $\alpha = 1/u_u^2$ and $\beta = 1/(u_p u_u^2)$. A well-known fact from the statistics literature is that $\text{E}(B) = \alpha/\beta$ and $\text{Var}(B) = \alpha/\beta^2$, hence

$$\begin{aligned} \text{E}(B) &= \frac{1}{u_u^2} u_p u_u^2 &= u_p \\ \text{Var}(B) &= \frac{1}{u_u^2}\left(u_p u_u^2\right)^2 &= u_p^2 u_u^2 \end{aligned}$$

The price-volume relationship for a bidder refers to the relationship between the expected awarded impression volume and the bid price $u_p$ for a fixed value of bid uncertainty $u_u$. Figure 2 illustrates the concept of price-volume relationship for an artificially constructed competitive landscape. The bidder is interested in impressions from 22 different segments and the highest competing bid price on each segment is reflected by a step in the price-volume curve. The height of each step corresponds to the available inventory volume for the corresponding inventory segment. Actually constructing price-volume curves for bidders in a large and highly segmented network requires centralized information and is computationally expensive. It is therefore rarely done in practice. However, the nature of price-volume curves is fundamental and of great importance. The left plot shows the price-volume relationship when no Heisenberg bidding is used, and the right plot shows the price-volume relationship for two different non-zero values of $u_u$.

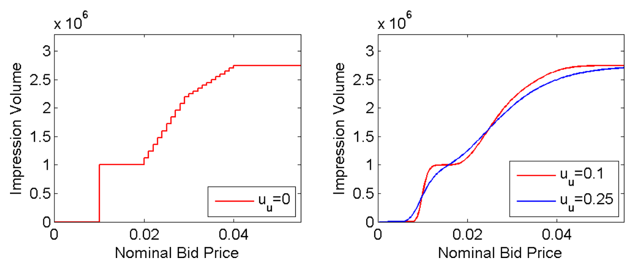


Fig. 2. Artificial price-volume curves for three different values of $u_u$.

A significant value of Heisenberg bidding is that $u_u > 0$ renders the otherwise discontinuous price-volume curve smooth. This opens up for the possibility to predict the volume locally around some operating point and enables an algorithm designer to make use of extensive mathematical results available only for smooth systems. It follows that Heisenberg bidding may enable efficient non-cooperative bidding strategies that enhance learning performance and scalability. Of greatest relevance to this paper and exploited in Section VI is that $u_u > 0$ reduces the effective plant gain at steps and increases the effective plant gain along plateaus, where the plant gain otherwise would be $\infty$ and 0, respectively.

## IV. Problem Formulation

Consider a network of bidders competing over impressions on Internet. The impressions represent showing of e.g. display ads or content promotions. Each bidder may be interested only in some segments of the available impressions, where an impression's *segment* is defined by user characteristics (age, gender, geographic location, interests, frequency, etc.), web site and content properties (site-slot location and content category), and/or time. Moreover, a bidder may submit different bids for impressions from different segments.

The control objective is to determine the proper bids for one of the bidders (referred to as *our* bidder) such that the number of impressions awarded to this bidder tracks a desired reference rate $\bar{u}_c(t)$. The control algorithm must handle noisy data related to e.g. stochastic traffic, and a dynamic competitive landscape related to competing bidders entering or leaving the network and competing bids changing over time. Last but not least, the control algorithm must handle seasonal patterns in traffic and discontinuous relationships between bid price and awarded impression volume. The reference signal $\bar{u}_c(t)$ is defined without taking supply seasonality into account, but the controller is expected to distribute the awarded number of impressions according to the seasonality.

## V. Plant Model

The plant as perceived by each bidder has the control input signals $u_p(t)$ and $u_u(t)$ representing the bid price and bid uncertainty used for the bidder at time $t$. The plant output signal is denoted $y(t)$ and represents a measured rate of awarded impressions $n_I(t-\Delta)/\Delta$, where $n_I(t)$ is the number of impressions awarded in the interval $[t, t+\Delta]$, where $\Delta$ is the sampling time.

The exogenous input signals consists of all competing bids, the arrival process of impression requests, and the reference signal. There may be a very large number of competing bids, but assume each bidder submits only one bid price and bid uncertainty per segment $j = 1, \ldots, n$.

Let $N_{tot,j}(t)$ denote the expected arrival of impression requests on segment $j$ in the interval time $[t, t+\Delta)$. Our bidder will be awarded some or all of $N_{tot,j}(t)$ depending on the relationship between $u_p(t)$, $u_u(t)$, competing bids, and the random draws in Heisenberg Perturbation.

If the highest competing bid price on segment $j$ at time $t$ is given by $\theta_j^\star(t)$, then the number of impressions awarded to our bidder on segment $j$ is a Binomial random number given by $n_{I,j}(t) \sim \text{Binomial}(N_{tot,j}(t), \text{Prob}(B > \theta_j^\star(t)))$.

Given a large value of $N_{tot,j}(t)$ the actual allocation $n_{I,j}(t)$ will be very close to the expected value $\text{E}\left(n_{I,j}(t)\right) = N_{tot,j}(t)\text{Prob}\left(B > \theta_j^\star(t)\right)$. The total number of impressions $n_I(t)$ awarded to our bidder is the sum of impressions awarded in each segment; i.e., $n_I(t) = \sum_{j=1}^n n_{I,j}(t)$.

In applications such as online advertising and content distribution $N_{tot,j}(t)$ is stochastic and time-varying. An approximation that works well for the purpose of control design and that we adopt is that the base-level volume is lognormal, and that the dominant time-varying effect is a 24 hour seasonality. In particular, we assume.

$$N_{tot,j}(t) \quad = \quad N_{0,j}\left(1+\beta_1 \sin\left(\frac{2\pi t}{24}+\phi_1\right)+\beta_2 \sin\left(\frac{4\pi t}{24}+\phi_2\right)\right)e^{\epsilon(t)}, \quad (1)$$

where $\epsilon(t) \sim N(0, \sigma_j^2)$, and where the various parameters are chosen so that $N_{tot,j}(t) \geq 0$ for all $t$. The two harmonics in the above model may be interpreted as a truncated Fourier series expansion of some arbitrary 24 hour periodic function. In practice, two harmonics is considered capturing Internet traffic seasonality well.

Figure 3 provides an example of impression volume over

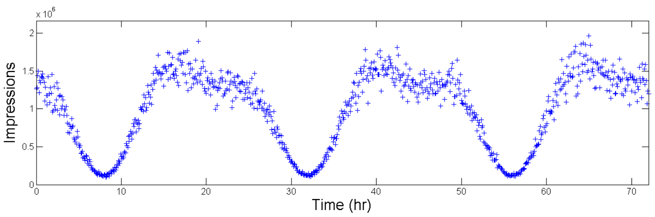


Fig. 3. Example of impression time series data generated according to (1) with model parameters $N_{0,j} = 10^6$, $\beta_1 = 0.63$, $\phi_1 = 2.76$, $\beta_2 = 0.26$, $\phi_2 = 0.39$, and $\sigma_j = 0.1$.

time. The example shows the effect of the two-harmonic seasonal model, and the stochastic impression volume.

## VI. CONTROL

In this section we illustrate how adaptive feedback control together with Heisenberg bidding can be used for decentralized control of a network of bidders. Without Heisenberg bidding, which is used to shape the plant gain, stability and robustness would typically require a very conservative and slow controller compromising the performance. Limited space prohibit us from discussing stability, robustness, and response to load or measurement noise.

### A. Preliminary Price Control Design

To better understand the ideal dynamics of the main control loop, consider first the simplified linear (but periodic) plant model defined by $y(t) = K_p(1+g(t))u_p(t-\Delta_p)$, where $K_p > 0$ is a constant plant gain, $\Delta_p > 0$ is the plant delay, and $g(t+T) = g(t) > -1$ for all $t$ and some $T$.

The objective is to regulate the plant so that $y(t)$ tracks $\bar{u}_c(t)$ on average sufficiently fast and with smallest possible variations in $u_p(t)$. To track $\bar{u}_c(t)$ on average means an error feedback controller does not need to use $\bar{u}_c(t)$ directly as the reference input signal, but may use a pre-filtered version of the signal given by $u_c(t) = (1+g(t))\bar{u}_c(t)$.

The error feedback controller is implemented by means of a PI-controller [7] defined by

$$u_p(t) \quad = \quad K_c\left(e(t) + \frac{1}{T_i}\int_0^t e(\tau)d\tau\right),$$

where $e(t) = u_c(t) - y(t)$.

The closed loop control system is illustrated as a block diagram in Figure 4.

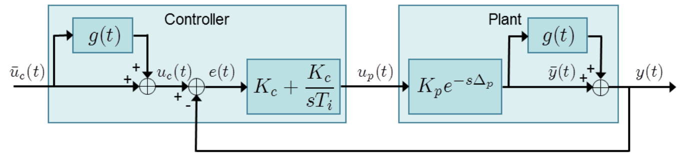


Fig. 4. Closed loop system with the simplified $T$-periodic linear plant.

A state space representation of the controller is given by

$$\frac{dx_1}{dt} = \frac{1}{T_i}e(t) \quad (2)$$
$$u_p(t) = K_c x_1(t) + K_c e(t). \quad (3)$$

Decompose the plant in a time-invariant and a time-varying component as

$$\bar{y}(t) = K_p u_p(t-\Delta_p)$$
$$y(t) = (1+g(t))\bar{y}(t).$$

Consider first the time-invariant component $\bar{y}(t) = K_p u_p(t-\Delta_p)$. Its Laplace transform is $\bar{Y}(s) = K_p e^{-s\Delta_p}U_p(s)$ and a first order Padé approximation of the transfer function $K_p e^{-s\Delta_p}$ is given by

$$K_p e^{-s\Delta_p} \approx K_p \frac{1-s\Delta_p/2}{1+s\Delta_p/2},$$

which can be expressed in observable canonical state-space form as

$$\frac{dx_2}{dt} = -\frac{2}{\Delta_p}x_2(t) + \frac{4K_p}{\Delta_p}u_p(t) \quad (4)$$
$$\bar{y}(t) = x_2(t) - K_p u_p(t). \quad (5)$$

Adjoining the dynamics of controller (2)-(3) and of the time-invariant part of the plant (4)-(5) yields

$$\frac{dx_1}{dt} = \frac{1}{T_i}e(t)$$
$$\frac{dx_2}{dt} = \frac{4K_cK_p}{\Delta_p}x_1(t) - \frac{2}{\Delta_p}x_2(t) + \frac{4K_cK_p}{\Delta_p}e(t)$$
$$\bar{y}(t) = -K_cK_p x_1(t) + x_2(t) - K_cK_p e(t)$$

This can be written in matrix form as

$$\frac{dx}{dt} = Ax(t) + Be(t)$$
$$\bar{y}(t) = Cx(t) + De(t)$$

where $A, B, C$, and $D$ are given by

$$A = \begin{bmatrix} 0 & 0 \\ \frac{4K_cK_p}{\Delta_p} & -\frac{2}{\Delta_p} \end{bmatrix}, \quad B = \begin{bmatrix} \frac{1}{T_i} \\ \frac{4K_cK_p}{\Delta_p} \end{bmatrix},$$
$$C = \begin{bmatrix} -K_cK_p & 1 \end{bmatrix}, \quad D = -K_cK_p.$$

Next introduce the time-varying parts of the system dynamics, close the loop, and consider $e(t) = u_c(t) - y(t)$ the output of interest.

$$e(t) = (1+g(t))\bar{u}_c(t) - (1+g(t))\bar{y}(t)$$
$$= (1+g(t))(\bar{u}_c(t) - Cx(t) + K_cK_p e(t))$$

Assume $K_c K_p(1+g(t)) \neq 1$. Then rearrange this equation so that $e(t)$ appears alone and only on the left hand side of the equation

$$\begin{aligned} e(t) &= \frac{1+g(t)}{1-K_cK_p(1+g(t))}(\bar{u}_c(t)-Cx(t)) \\ &= \breve{g}(t)(\bar{u}_c(t)-Cx(t)) \end{aligned}$$

where

$$\breve{g}(t) = \frac{1+g(t)}{1-K_cK_p(1+g(t))}$$

It follows that the closed-loop dynamics is described by

$$\begin{aligned} \frac{dx}{dt} &= Ax(t)+B\breve{g}(t)(\bar{u}_c(t)-Cx(t)) \\ &= (A-\breve{g}(t)BC)\,x(t)+\breve{g}(t)B\bar{u}_c(t) \\ e(t) &= \breve{g}(t)(\bar{u}_c(t)-Cx(t)) \\ &= -\breve{g}(t)Cx(t)+\breve{g}(t)\bar{u}_c(t) \end{aligned}$$

or equivalently, by

$$\frac{dx}{dt} = \tilde{A}(t)x(t)+\tilde{B}(t)\bar{u}_c(t) \quad (6)$$

$$e(t) = \tilde{C}(t)x(t)+\tilde{D}(t)\bar{u}_c(t) \quad (7)$$

where

$$\begin{aligned} \tilde{A}(t) &= A-\breve{g}(t)BC, & \tilde{B}(t) &= \breve{g}(t)B, \\ \tilde{C}(t) &= -\breve{g}(t)C, & \tilde{D}(t) &= \breve{g}(t). \end{aligned}$$

It is straight forward to show that (6)-(7) has the unique equilibrium solution

$$x^\star = \begin{bmatrix} \frac{1}{K_cK_p} \\ 2 \end{bmatrix}\bar{u}_c \quad \text{and} \quad e^\star = 0.$$

It is not immediately clear under what conditions the equilibrium solution is stable. For example, for linear time-varying systems, uniform asymptotic stability cannot be characterized by the location of the eigenvalues of $A(t)$ [8]; i.e., for stability it is not sufficient that the eigenvalues for all $t$ lie in the open left-half plane. The evaluation of stability and robustness is outside the scope of this paper.

### *B. Adaptive Heisenberg Bidding Based Design*

Now consider the realistic and discontinuities plant model discussed in Section V and assume $T_i$ is chosen much larger than $\Delta_p$. The dynamics of the controller then dominates the dynamics of the plant and we may approximate the plant by a static relationship (set $\Delta_p = 0$). It follows that the input-output relationship of the plant is given by

$$y(t) = (1+g(t))f(u_p(t),u_u(t)), \quad (8)$$

where $g(t)$ is a $T$-periodic and slowly varying function satisfying $g(t) > -1$ and $T \gg T_i$; and where the properties of $f(u_p,u_u)$ is discussed in detail in section V (and illustrated in Figure 2/Section III).

Let us again consider a standard PI-controller for the primary control loop.

$$\frac{dx}{dt} = \frac{1}{T_i}e(t) \quad (9)$$

$$u_p(t) = K_cx(t)+K_ce(t) \quad (10)$$

The plant model (8) can be linearized around any operating point $u_p$ and $u_u > 0$. Treating the intercept term in the linearization as a load disturbance (and ignoring it), the loop transfer dynamics is

$$\begin{aligned} y(t) &= (1+g(t))K_pu_p(t) \\ &= (1+g(t))K_cK_p(x(t)+e(t)) \end{aligned}$$

Proceeding as in Section VI-A (but under the assumption $\Delta_p = 0$), the control error is given by

$$\begin{aligned} e(t) &= (1+g(t))\bar{u}_c(t)-y(t) \\ &= (1+g(t))\,(\bar{u}_c(t)-K_cK_px(t)-K_cK_pe(t)) \end{aligned} \quad (11)$$

Isolating $e(t)$ on the left-hand side yields

$$\begin{aligned} e(t) &= \frac{K_cK_p(1+g(t))}{1+K_cK_p(1+g(t))}\left(-x(t)+\frac{1}{K_cK_p}\bar{u}_c(t)\right) \\ &= \tilde{g}(t)\left(-x(t)+\frac{1}{K_cK_p}\bar{u}_c(t)\right) \end{aligned}$$

where

$$\tilde{g}(t) = \frac{K_cK_p(1+g(t))}{1+K_cK_p(1+g(t))} \quad (12)$$

Combining the controller dynamics (9)-(10) and the expression for the control error (11) results in the closed loop dynamics

$$\begin{aligned} \frac{dx}{dt} &= \frac{\tilde{g}(t)}{T_i}\left(-x(t)+\frac{1}{K_cK_p}\bar{u}_c(t)\right) \\ y(t) &= K_cK_p(1+g(t))\left((1-\tilde{g}(t))\,x(t)+\frac{\tilde{g}(t)}{K_cK_p}\bar{u}_c(t)\right) \end{aligned}$$

where $g(t+T) = g(t)$ and $\tilde{g}(t+T) = \tilde{g}(t)$. We obtain the following linear periodic system

$$\frac{dx}{dt} = A(t)x(t)+B(t)\bar{u}_c(t) \quad (13)$$

$$y(t) = C(t)x(t)+D(t)\bar{u}_c(t) \quad (14)$$

where

$$\begin{aligned} A(t) &= -\frac{\tilde{g}(t)}{T_i} \\ B(t) &= \frac{\tilde{g}(t)}{T_iK_cK_p} \\ C(t) &= K_cK_p(1+g(t))\,(1-\tilde{g}(t)) \\ D(t) &= (1+g(t))\tilde{g}(t) \end{aligned}$$

The problem now is that $K_p$ is unknown, dependent on $u_p$ and $u_u$, and may vary over time as the competitive landscape changes. We opt to handle this by means of a model-reference adaptive controller implemented using the

MIT rule [9], where the adaptive controller adjusts $K_c$. The desired closed-loop response is specified by a model whose output is $y_m$. Furthermore, $e_m$ is the error between the output $y$ of the closed-loop system and the desired closed-loop response $y_m$; i.e., $e_m(t) = y(t) - y_m(t)$.

With the MIT rule $K_c$ is adjusted gradually to minimize the loss function $e_m^2/2$. This is achieved by incremental changes of $K_c$ along the negative gradient of the loss function; i.e., $dK_c/dt = -\gamma e_m \partial e_m/\partial K_c$ for some carefully chosen value of $\gamma$.

Suppose next that the controller by means of proper choices of $T_i$ and $K_c$ is much faster than $A$, $B$, $C$, and $D$. We can then approximate the dynamics at each time point $t$ (in the near future) as time-invariant and transform the closed loop dynamics into Laplace domain pretending $A$, $B$, $C$, and $D$ are independent of time; i.e.,

$$\begin{aligned} Y(s) &= \left(C(t)(s-A(t))^{-1}B(t) + D(t)\right)\bar{u}_c(s) \\ &= \left(\frac{K_cK_p(1+g(t))(1-\tilde{g}(t))\tilde{g}(t)}{(s+\tilde{g}(t)/T_i)T_iK_cK_p} + (1+g(t))\tilde{g}(t)\right)\bar{U}_c(s) \\ &= (1+g(t))\tilde{g}(t)\left(\frac{1-\tilde{g}(t)}{sT_i+\tilde{g}(t)} + 1\right)\bar{U}_c(s) \end{aligned}$$

Using (12) we can rewrite this as

$$Y(s) = \frac{K_cK_p(1+g(t))^2(sT_i+1)}{sT_i(1+K_cK_p(1+g(t))) + K_cK_p(1+g(t))}\bar{U}_c(s)$$

Let the desired dynamics be defined by

$$Y_m(s) = \frac{K_m(1+g(t))^2(sT_i+1)}{sT_i(1+K_m(1+g(t))) + K_m(1+g(t))}\bar{U}_c(s) \tag{15}$$

where $K_m$ is a design parameter. By adjusting the controller gain $K_c$ such that $K_cK_p \approx K_m$, then $Y(s) \approx Y_m(s)$.

The adaptation error is defined by $E_m(s) = Y(s) - Y_m(s)$, and the partial derivative of $E_m(s)$ used in the adaptation can easily be shown to be

$$\frac{\partial E_m}{\partial K_c} = \frac{K_p(1+g(t))^2(sT_i+1)sT_i}{(sT_i(1+K_cK_p(1+g(t))) + K_cK_p(1+g(t)))^2}\bar{U}_c(s)$$

This formula cannot be used directly since $K_p$ is unknown. An approximation is therefore necessary. One such approximation that we shall use is $K_cK_p \approx K_m$. Then $K_p \approx K_m/K_c$ and we have $\partial E_m/\partial K_c = Z(s)/K_c$, where

$$\begin{aligned} Z(s) &= \frac{\tilde{g}_m(t)^2}{K_m}\frac{(sT_i+1)sT_i}{(sT_i+\tilde{g}_m(t))^2}\bar{U}_c(s) \\ \tilde{g}_m(t) &= \frac{K_m(1+g(t))}{1+K_m(1+g(t))} \end{aligned}$$

The MIT rule can now be expressed as

$$\frac{dK_c}{dt} = -\gamma e_m(t)\frac{z(t)}{K_c(t)}$$

Before we assemble the derived formulas for updating $K_c$, let us propose how to update the uncertainty bid $u_u(t)$. The idea is to shape the plant gain as needed, and the need is present near discontinuities in the price-volume curve to ensure there exists a solution to the control problem, and along plateaus ($K_p \approx 0$) to avoid the adaptive controller ramping up the controller gain $K_c$ to dangerous values that would be damaging if the plateau suddenly ends with a sudden step.

A simple scheme is to establish a one-to-one relationship $G_u$ between $u_u(t)$ and $K_c(t)$ defined by design parameters $u_u^0$ and $K_c^0$ as follows: $u_u(t) = G_u(K_c(t))$, where $G_u(K_c(t)) = u_u^0 + K_u\left(K_c(t) - K_c^0\right)^2$.

The results so far are illustrated in Figure 5.

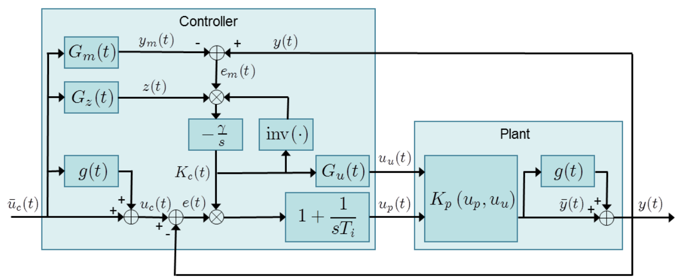


Fig. 5. Block diagram of the derived adaptive controller implemented in the closed loop system.

To implement each module of the block diagram we need state-space representations of $G_m : \bar{u}_c \to y_m$ and $G_z : \bar{u}_c \to z$. It is straight forward (but somewhat tedious) to obtain

$$\begin{aligned} \frac{dx_m}{dt} &= -\frac{\tilde{g}_m(t)}{T_i}x_m(t) + \bar{u}_c(t) \\ y_m(t) &= \frac{(1+g(t))\tilde{g}_m(t)(1-\tilde{g}_m(t))}{T_i}x_m(t) \\ &\quad +(1+g(t))\tilde{g}_m(t)\bar{u}_c(t) \end{aligned}$$

and

$$\begin{aligned} \frac{dx_z}{dt} &= \begin{bmatrix} -\frac{2\tilde{g}_m(t)}{T_i} & -\frac{\tilde{g}_m(t)^2}{T_i^2} \\ 1 & 0 \end{bmatrix} x_z(t) + \begin{bmatrix} 1 \\ 0 \end{bmatrix}\bar{u}_c(t) \\ z(t) &= \frac{\tilde{g}_m(t)^2}{K_m}\begin{bmatrix} \frac{1-2T_i\tilde{g}_m(t)}{T_i} & -\frac{\tilde{g}_m(t)^2}{T_i^2} \end{bmatrix} x_z(t) + \frac{\tilde{g}_m(t)^2}{K_m}\bar{u}_c(t) \end{aligned}$$

The continuous-time adaptive controller must be discretized in practical implementations, but given the above state-space representations that is a trivial exercise.

## VII. SIMULATION RESULTS

A slight variation of the algorithm developed in this paper is implemented and is being rolled out on `www.aol.com`, a major web portal, as part of the proprietary ContentLearn™ system [10]. It is used to satisfy sponsorship commitments where Aol has agreed to drive a certain amount of traffic to specific sites hosting advertisement. This is accomplished by promoting content (not advertisement) from the sponsored site in just the right amount on the high traffic home portal. By means of feedback control the content promotions are shown to the number of users that is required for the traffic to the sponsored landing page to reach the commitment. The goal in this configuration is to pace the click stream, not impression stream as assumed in this paper. However, it is straight-forward to generalize the algorithm in Section VI to the case of click volume control.

To appreciate the basic behavior of Heisenberg bidding based control we limit our testing in this paper on simulation results. Suppose the expected daily available impression volume and the competitive landscape is given by the price-volume curve in Figure 2. The competitive landscape is assumed to be static throughout the scenario, but the supply of impressions is stochastic and obeys the seasonality model

$$g(t) = 0.63 \sin\left(\frac{2\pi t}{24} + 2.76\right) + 0.26 \sin\left(\frac{4\pi t}{24} + 0.39\right),$$

which is a good approximation of Internet traffic in the US. Assume the hourly rate of available impression volume at time $t$ is $N_{tot,j}(t) = N_{0,j}^{daily}\left(1 + g(t)\right)e^{\epsilon(t)}$ where $N_{0,j}^{daily}$ is the expected daily available impression volume in segment $j$, and $\epsilon(t) \sim N(0, 0.1^2)$.

Suppose our bidder is subject to the daily delivery goal:

$$\bar{u}_c^{daily}(t) = \begin{cases} 450,000 & \text{if } 0 \le t < 48 \\ 2,450,000 & \text{if } 48 \le t < 96 \\ 1,450,000 & \text{if } 96 \le t \le 144 \end{cases},$$

which implies that the reference signal $\bar{u}_c(t) = \bar{u}_c^{daily}(t)/24$.

The controller derived in Section VI-B and illustrated in Figure 5 is now used to solve this problem, but with three modification. First, the output $y$ is filtered as follows before being fed to the controller

$$Y_f(s) = \frac{1}{1 + sT_f + (sT_f)^2/2} Y(s)$$

Next, the plant and controller are discretized with a sampling time of $\Delta = 1/12$ hours, and finally, the plant is subject to a delay of $\Delta_p = \Delta$.

The following parameter values are used: $T_f = 0.0938$, $T_i = 0.3750$, $K_m = 0.5$, $\gamma = 10^{-21}$, $K_c(0) = 8 \cdot 10^{-7}$, $u_p(0) = 0.005$, $u_u^0 = 0.1$, $K_u = 10^5$, and $K_c^0 = 8 \cdot 10^{-7}$.

The closed loop result is shown in Figure 6. The red

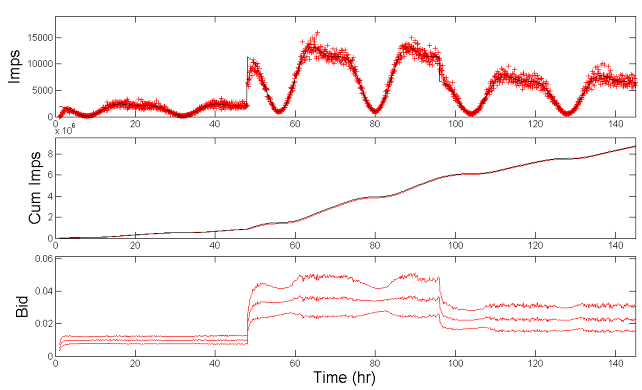


Fig. 6. Closed loop result demonstrating the step response of the proposed controller. The top and middle plots show the marginal and cumulative tracking performance, and the bottom plot shows $u_p$ (center curve) and a 95% bid interval defined by the bid distribution $B \sim \text{Gamma}(1/(u_u^2), 1/(u_p u_u^2))$.

markers in the top plot display the number of awarded impressions in each sampling interval while the black curve shows the desired number of impressions per sample; i.e., $u_c(t)\Delta = \bar{u}_c(t)(1 + g(t))\Delta$. The middle plot shows the information in the first plot cumulatively. The bottom plot encodes bid price $u_p$ and bid uncertainty $u_u$. The middle curve shows $u_p$ while the band defined by the outer two curves define a 95% bid interval defined by the bid distribution $B \sim \text{Gamma}\left(1/(u_u^2), 1/(u_p u_u^2)\right)$.

Note how rapidly the controller responds to step changes in the reference signal and how robustly it handles seasonality and the stochastic element of the Internet traffic.

## VIII. Concluding Remarks

We have proposed an approach to control design for agents in a network of bidders, and illustrated the approach with a model-reference adaptive controller designed to regulate Internet traffic. The underlying approach is based on bid randomization to make efficient and robust control of real-world auction networks possible. Indeed, solutions available in the literature are based on assumptions that are invalid in practice, and violations of these assumptions typically destroy possible optimality (sometimes dramatically).

An added benefit of the proposed bid randomization is that it turns the system tractable for mathematical analysis using techniques only available for smooth systems. This in itself is a major advantage since it expands on the set of tools that can be applied in future control designs.

The results of this paper are only a first step in the utilization of the proposed bid randomization. Many interesting topics for further studies exist. They relate to designing other controllers, analysing robustness and sensitivity, establishing conditions for stability, determining criteria indicating the onset of limit cycles or chaos.

## References


[1] Niklas Karlsson. *Computerized Systems and Method related to Controlled Content Optimization (U.S. Provisional Application No. 61/882,845)*. USPTO, 2013.

[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, 2013.

[3] Deepak Agarwal, Bee-Chung Chen, Pradheep Elango, Nitin Motgi, Seung-Taek Park, Raghu Ramakrishnan, Scott Roy, and Joe Zachariah. *Online Models for Content Optimization*. Proceeedings of the 2008 NIPS conference, Vancouver, Canada, Dec 8 - 11, 2008.

[4] Deepak Agarwal, Bee-Chung Chen, Pradheep Elango, and Raghu Ramakrishnan. Content recommendation on web portals. *Commun. ACM*, 56(6):92–101, 2013.

[5] Niklas Karlsson. *Systems and Methods For Advertisement And Content Optimization Over An Electronic Network, United States Patent Application 13/734,587*. USPTO, 2013.

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

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

[8] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, second edition, 1996.

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

[10] Niklas Karlsson, Jianlong Zhang, Robert Luenberger, Scott Robert Strickland, Yih-Shin Cindy Huang, Seyed Mohammad Ziaemohseni, Li Yang, and Hans W. Uhlig. *Systems and Methods For Displaying Digital Content And Advertisements Over Electronic Networks, United States Patent Application 20130197994*. USPTO, 2012.