Electronic Commerce Research and Applications 11 (2012) 537–547

Contents lists available at SciVerse ScienceDirect

Electronic Commerce Research and Applications

journal homepage: www.elsevier.com/locate/ecra

Co-evolution-based mechanism design for sponsored search advertising

Yong Yuan [a,b], Daniel Zeng [a,c,*]

[a] State Key Laboratory of Management and Control for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing, China
[b] Beijing Engineering Research Center of Intelligent Systems and Technology, Beijing, China
[c] Department of Management Information Systems, University of Arizona, Tucson, Arizona, USA

ARTICLE INFO

Article history:
Available online 25 April 2012

Keywords:
Sponsored search advertising
Mechanism design
Co-evolutionary simulation
Niche-based evolution

ABSTRACT

Sponsored search advertising (SSA), the primary revenue source of Web search engine companies, has become the dominant form of online advertising. Search engine companies, such as Google and Baidu, are naturally interested in SSA mechanism design with the aim to improve the overall effectiveness and profitability of SSA ecosystems. Due to model intractability, however, traditional game theory and mechanism design frameworks provide only limited help as to the design and evaluation of practical SSA mechanisms. In this paper, we propose a niche-based co-evolutionary simulation approach, aiming at computationally evaluating SSA auction mechanisms based on advertisers' equilibrium bidding behavior generated through co-evolution of their bidding strategies. Using this approach, we evaluate and compare key performance measures of several practical SSA auction mechanisms, including the generalized first and second price auction, the Vickrey–Clarke–Groves mechanism, and a novel hybrid mechanism adopted by sogou.com, a major search engine in China.

1. Introduction

In sponsored search advertising (SSA), online advertisers bid for keyword-specific advertisements to appear alongside the organic search results on Web search result pages. With the promise of precise and in-context customer targeting, SSA provides an effective way of monetizing Web search queries. Within a decade, it has evolved into the dominant form of online advertising and becomes an industry on its own. In 2010, SSA constituted the largest category share (46%) of the $26 billion online advertisement spending in the U.S. markets, far exceeding the share of display advertisement, the second largest category, 24%.[1] SSA is also the primary revenue source of Web search engine companies. In recent years, it accounted for more than 96% and 99.9% of Google and Baidu's international revenues, respectively.[2]

The basic economic institution behind most SSA platforms, such as Google's AdWords and Baidu's Phoenix Nest, is keyword-based position auction. In this type of auction, advertisers selling similar products or services bid for the same keywords on an SSA platform. Once a relevant search query arrives, an auction will be conducted to determine the rank position and the associated payment of winning advertisements. When a user clicks on an advertisement, she will be sent to the landing page on the website of the corresponding advertiser, who in turn pays the search engine.

From search engine companies' point of view, SSA auction mechanism design, i.e., the choice of auction mechanism or format with the ranking and pricing schemes at its core, is of paramount importance. Different auction mechanisms induce different types of advertiser bidding behavior, which in turn determine advertisers' revenues and search engine companies' profitability. In the long run, the stability and sustainability of the SSA ecosystem, to a large degree, hinges on SSA auction mechanism design as well.

A variety of auction mechanisms have emerged during the evolution of SSA since its appearance in 1998. The pioneer of SSA, GoTo.com (then Overture, now part of Yahoo!), used the generalized first-price (GFP) auction, in which advertisers were ranked by and pay their own bids. In 2002, Google started to use the generalized second-price (GSP) auction, which ranks advertisers by their bids but charges them, if their advertisements are clicked by Web users, by the next highest bids. Currently, most major search engines around the world have adopted a variant of the GSP mechanism, in which advertisers are ranked by the product of their own bids and search engine-assigned quality scores. In the meanwhile, other auction formats have been experimented and used as well. A prominent example is a hybrid auction format adopted by sogou.com, China's third largest search engine. In this auction, advertisers are ranked by their bids but allowed to select to pay following either the first or second pricing scheme. In this

* Corresponding author at: Department of Management Information Systems, University of Arizona, Tucson, Arizona, USA.

E-mail addresses: yong.yuan@ia.ac.cn (Y. Yuan), zeng@email.arizona.edu (D. Zeng).

[1] http://www.iab.net/insights_research/947883/adrevenuereport#tabs-10.

[2] http://investor.google.com/financial/tables.html, and http://ir.baidu.com/phoenix.zhtml?c=188488&p=irol-reportsAnnual.

1567-4223/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.elerap.2012.03.002

paper, we refer to this auction format as the First–Second-Price (FSP) auction.

In the literature, the Vickrey–Clarke–Groves (VCG) and laddered auctions have also been studied (Aggarwal et al. 2006). These auction formats offer nice properties such as incentive compatibility, typically aimed at optimizing the profitability of search engines, and at reducing the possibility of advertisers trying to "game the system." As a result, market efficiency can be achieved, i.e., advertisers with higher per-click values are more likely to win better advertisement slots.

Due to its practical significance, SSA auction mechanism design has attracted a lot of attention from the research community in recent years. Most existing work evaluates auction mechanisms using frameworks and tools from game theory and, in particular, mechanism design theory. As SSA platforms evolve and associated auction rules undergo changes in dynamic online business settings, formal mathematical analysis based on mechanism design theory quickly becomes inadequate in dealing with design and evaluation of practical SSA auction mechanisms. It has become a major challenge for search engines to understand advertisers' bidding behavior and assess the outcome of complex SSA auctions. As a result, these companies can only rely on ad hoc anecdotal evidence or limited scenario-based comparisons to make major decisions concerning auction rules, increasing the risk and hindering innovation in the SSA space.

From a research perspective, with the end goal of helping the SSA ecosystem maintain its overall effectiveness, profitability, and stability, there is a critical need to develop new research methodologies for SSA mechanism design. However, to the best of our knowledge, research in this area is severely lacking. Our research is targeted at filling in this important gap. We propose a niche-based co-evolutionary approach, aiming at automatically searching for equilibrium bidding behavior of rational advertisers in SSA auctions, and in turn evaluating and designing alternative SSA auction mechanisms. Niche plays a central role in evolutionary divergence during the ecological speciation process, and can be used in co-evolutionary simulations to construct the equilibrium continuum of SSA auctions by forming and maintaining stable sub-populations, each converging at a single equilibrium. Compared with the analytical mechanism design approach, co-evolutionary simulation offers the following advantages. First, it evaluates the macro-scope properties of SSA auctions by simulating long-term co-evolution and co-adaptation of advertisers' micro-scope bidding behavior. As such, the simulation process is robust to input mechanisms, and can be adapted to assess various kinds of SSA mechanisms with only minor modifications. Second, co-evolutionary simulation can be used to optimize advertisers' bidding strategies via searching through the entire strategy space. This can help search engines better understand advertisers' behavior in a specific SSA mechanism.

This paper makes the following contributions. Methodologically, the reported work is the first attempt to apply the co-evolutionary simulation approach to auction mechanism design in SSA contexts. We also propose a novel niche-based co-evolutionary algorithm to help design and evaluate SSA auction mechanisms. Practically, our research can help Web search engines better understand advertisers' complex bidding dynamics in SSA auctions, and computationally evaluate SSA auctions' performance. From the perspective of competing advertisers, our research and algorithm can help them discover all kinds of possible equilibrium bidding strategies, and analyze their marketing effectiveness. As a result, a variety of observed bidding behavior to "game the system" (Zhou and Lukose 2006), as well as irrational or even malicious bidding behavior (Iyengar et al. 2007), may be recognized and reduced in SSA practice.

The remainder of this paper is organized as follows. Section 2 provides a brief review of the SSA auction mechanism design literature. In Section 3, we discuss performance measures concerning SSA auction mechanisms and the rationale behind co-evolutionary auction simulation. We then present in detail our proposed niche-based co-evolutionary mechanism design. Several typical SSA auction mechanisms are evaluated through co-evolutionary simulations in Section 4. In Section 5, we summarize our research findings and discuss future research possibilities.

2. Literature review

Mechanism design has long been an active topic in auction research. Here we are mainly concerned with the mechanism design work in the SSA context. Two research streams, analytical and computational mechanism design, have been developed in the literature. Below we present a brief survey of these two lines of thoughts.

2.1. Analytical mechanism design

The classical mechanism design framework has been used to formally characterize the key properties of various SSA auction mechanisms. For instance, the early GFP mechanism has been proved to be unstable with price cycles in bids (Zhang and Feng 2005). In contrast, the prevailing GSP mechanism has a symmetric Nash equilibrium (SNE) continuum, and the bids of revenue-maximizing advertisers will converge to the lowest-price Nash equilibrium (LPNE) in SNE (Edelman et al. 2007). However, GSP is not a truthful mechanism, and its Nash equilibria (NE) beyond the SNE continuum is not yet fully explored. For theoretical analysis, researchers have investigated truthful SSA mechanisms such as the VCG and laddered auctions (Aggarwal et al. 2006), in which advertisers are motivated to truthfully bid their private per-click values. Moreover, an optimal mechanism has been proposed to maximize search engines' expected revenue while achieving Bayesian incentive compatibility and individual rationality of advertisers (Garg and Narahari 2009).

Recent analytical research focuses mainly on mechanism design for emerging SSA formats. For instance, an execution-contingent VCG mechanism was proposed for SSA platforms operating on federated search engines (Ceppi et al. 2011). Multi-slot SSA auctions have been studied, in which an advertiser can bid for multiple advertisement slots simultaneously (Deng et al. 2010). Several extended forms of advertisers' utility functions, e.g., a linear form with identical slopes and a single discontinuity, a piece-wise linear form with non-identical slopes and multi-discontinuities, have also been developed to improve the expressiveness of SSA mechanisms (Aggarwal et al. 2009, Duting et al. 2011).

In general, analytical mechanism design has the following limitations. First, due to inherent analytical complexity, it is difficult to mathematically evaluate complex SSA mechanisms in dynamic online environments. Second, the analytical approach is usually highly sensitive to auction mechanisms. Even a minor modification to the mechanism can lead to totally different analyses and solutions. For instance, if advertisers aim to maximize their revenues and the rivals' payments in competitive SSA markets, their bids in a GSP auction will converge to the highest-price equilibrium in the SNE continuum, instead of the LPNE (Yuan et al. 2011). Third, the standard game-theoretic analysis cannot reveal such dynamic properties as stability and robustness of equilibrium bidding strategies, and thus sheds limited light on which kind of equilibrium outcome is more likely to be observed in SSA auctions with a specific mechanism over the long run. These limitations have motivated research in computational mechanism design.

2.2. Computational mechanism design

Various kinds of computation-intensive techniques, such as experimental simulation, machine learning, data mining and evolutionary search, have been used in computational mechanism design research to better understand advertisers' bidding behavior in SSA auctions.

In the literature, Feng et al. (2007) compared the steady-state performance of four alternative SSA mechanisms, including the rank-by-bid, rank-by-revenue, rank-by-click through rate (CTR), and posted-price mechanisms, with systematic computational simulations. Balcan et al. (2005) proposed to formulate the problem of design of revenue-maximizing incentive-compatible SSA mechanisms as an algorithmic pricing problem with sample-complexity techniques from machine learning theory. Data mining techniques have been applied to inform empirical mechanism design. For instance, Ciaramita et al. (2008) presented a framework for learning and evaluating the SSA ranking systems based exclusively on click-data. Pardoe et al. (2010) proposed a data-driven adaptive methodology for SSA mechanism design, which incorporates prior general knowledge of bidding behavior to improve the effectiveness of mechanism design in dynamic and uncertain environment. Munsey et al. (2010) is the first to use evolutionary search in the SSA context, aiming to optimize advertisers' bids with historical data based genetic algorithms.

In this paper, we propose to apply the co-evolutionary simulation techniques in SSA mechanism design. This approach offers two advantages. First, as an explicit dynamic process, co-evolutionary simulation can automatically locate the equilibrium continuum of SSA auctions, and characterize the long-term dynamics and stability of specific behavior, strategies and mechanisms evolving over time. Second, the co-evolution process is largely independent of the auction mechanisms being evaluated, and can evolve the equilibrium outcomes adaptively.

Co-evolutionary mechanism design is not new in computational economics, and has been used successfully to design pricing rules for double auctions in a wholesale electricity marketplace (Phelps et al. 2002). In this paper, we make the first attempt to design and evaluate SSA mechanisms by co-evolutionary simulations. We also contribute to the methodology by introducing the "niche" technique, which can effectively prevent the co-evolution process from converging at a single equilibrium.

3. Co-evolutionary mechanism design for SSA auctions

In this section, we present the detailed methodology and related algorithms of our niche-based co-evolutionary mechanism design approach for SSA auctions. Auction mechanism design is concerned with designing auction mechanisms to meet specific performance requirements against self-interested bidders. Properly defining performance measures is a prerequisite for evaluating alternative mechanisms. We begin by presenting a simple SSA model and defining these performance measures in Section 3.1. We proceed in Section 3.2 by discussing the rationale behind co-evolutionary simulation, which serves as the base for SSA mechanism evaluation. In Section 3.3, we present the details of our proposed niche-based co-evolutionary mechanism design approach and the related simulation algorithm. We conclude this section by analyzing the computational complexity of this algorithm in Section 3.4.

3.1. An SSA model and performance measures

We consider a general SSA model with N revenue-maximizing advertisers competing for K slots on a specific keyword. The CTR of the k^{th} highest slot is denoted by x_k. We assume higher-placed slots have higher probabilities of being clicked, and thus $x_1 \geqslant x_2 \geqslant \cdots$. Each advertiser in slot $k \in [1,N]$ assigns a click with a private value v_k representing her maximum willingness to pay. Without loss of generality, we assume $v_1 \geqslant v_2 \geqslant \cdots$. The bids and payments of all advertisers are denoted as $b = \{b_1, b_2, \ldots\}$ and $p = \{p_1, p_2, \ldots\}$, respectively.

Based on this model, we define the following measures used to evaluate the performance of SSA auction mechanisms.

3.1.1. Market efficiency (social welfare)

This measure equals in value to the aggregated revenue of the SSA market participants, including the search engine and advertisers (Vijay and Perry 1997). In SSA auctions, an advertiser winning the k^{th} highest slot assigns a click with value v_k and pays p_k to the search engine when her advertisement is clicked. Advertisers winning no slots receive no clicks and thus do not pay. As such, an advertiser in slot k can expect to make a revenue of $(v_k - p_k)x_k$, and the total revenue of all advertisers is $\sum_{k=1}^{K}(v_k - p_k)x_k$. Analogously, the revenue received by the search engine equals the aggregated payments of all advertisers, i.e., $\sum_{k=1}^{K} p_k x_k$. Consequently, market efficiency is given by

$$ME = \sum_{k=1}^{K}(v_k - p_k)x_k + \sum_{k=1}^{K} p_k x_k = \sum_{k=1}^{K} v_k x_k$$

Market efficiency will be maximized (minimized) when advertisers are ranked in decreasing (increasing) order by their per-click values. However, a particular SSA auction may not be able to reach the maximum or minimum market efficiency. Using the niche-based co-evolutionary simulation framework, we can discover a large number of equilibria in the joint strategy space of advertisers engaging in SSA auctions. Assume we have found Θ equilibria, and the market efficiency realized in each equilibrium is denoted as $ME(i)$, $i \in [1, \Theta]$. Then we can define ME_{max}, ME_{min}, and ME_{avg} as the maximum, minimum and average market efficiency in the resulting equilibrium continuum, or formally

$$ME_{max} = \max\{ME(i)\}, \quad ME_{min} = \min\{ME(i)\}, \quad \text{and} \quad ME_{avg} = \frac{\sum_{i=1}^{\Theta} ME(i)}{\Theta}, \quad \text{where} \quad i \in [1, \Theta]$$

3.1.2. Revenue ratio

In SSA auctions, the search engine plays a non-cooperative game with advertisers to compete for auction surplus, or social welfare, equivalently. In order to characterize the revenue properties of SSA mechanisms, we define RR_{max}, RR_{min} and RR_{avg} as the maximum, minimum and average proportion of advertisers' aggregated revenues in the total auction surplus among all equilibria, where

$$RR = \frac{\sum_{k=1}^{K}(v_k - p_k)x_k}{\sum_{k=1}^{K} v_k x_k}$$

RR_{max}, RR_{min}, and RR_{avg} can be defined analogously to market efficiency. Obviously, search engines prefer SSA auction mechanisms with smaller revenue ratio, while advertisers prefer mechanisms with larger one.

3.1.3. Incentive compatibility

An SSA auction is said to be incentive compatible if all advertisers maximize their revenue when they truthfully bid their private per-click values (Dash et al. 2003). In co-evolutionary simulation, we consider an SSA mechanism to be incentive compatible if advertisers' bids always converge to their per-click values.

3.1.4. Output truthfulness

A ranking of advertisers is output truthful if it satisfies $v_t > v_k$ for $t < k$, $t, k \in [1, K]$ (Bu et al. 2010). We define *OT* as the proportion of the output-truthful equilibria among all equilibria found by co-evolutionary simulations.

3.1.5. Evolutionary stability

An SSA auction mechanism is evolutionarily stable if co-evolutionary simulations always converge stably to a unique equilibrium (Weibull 1995). Note that evolutionarily stable mechanisms will result in predictable bidding behavior and auction outcomes. We use static population variance and evolutionary mobile variance to characterize the co-evolutionary bidding dynamics and stability. The detailed definition of these two measures will be given in Section 3.3.

3.2. Co-evolutionary auction simulation

Inspired by co-adaptation and co-evolution of biological populations in natural ecosystems, co-evolutionary simulation simultaneously evolves multiple populations with coupled fitness. This approach has been successfully applied to strategy learning and optimization in games involving self-interested agents. Co-evolutionary simulation is particularly useful in SSA mechanism design for the following reasons. First, understanding equilibrium bidding strategies by advertisers plays an essential role in the evaluation of auction mechanisms. Co-evolutionary simulation has the potential to identify these stable and equilibrium strategies adaptively through the genetic search in advertisers' joint strategy space. Theoretically, it has been proved that in strategic environments, including SSA auctions, the co-evolutionary simulation process will converge at a globally optimal equilibrium with evolutionary stability and dynamic attainability (Apaloo et al. 2005). Second, co-evolutionary simulation has been applied in various kinds of auction mechanism design contexts with success. For instance, Cliff (2003) used co-evolutionary genetic algorithm to optimize parameter values for trading agents in online auction e-marketplaces. Phelps et al. (2002) applied co-evolutionary genetic programming and developed an auction pricing rule for double auctions in a wholesale electricity marketplace. These works provide useful modeling insights to tackle SSA mechanism design challenges. Third, the co-evolutionary simulation approach is largely insensitive to the auction mechanisms being evaluated. Investigating different kinds of SSA mechanisms only incurs relatively minor changes in the co-evolutionary simulation algorithm. Fourth, co-evolutionary simulation can serve as a foundation of a computable and implementable framework for SSA mechanism design. This famework may be used as a practical mechanism design and evaluation tool for Web search engines. These advantages have motivated us to apply co-evolutionary simulation to generate advertisers' equilibrium bidding behavior and in turn evaluate diverse SSA mechanisms.

Due to the non-cooperative nature of SSA auctions, we use the competitive co-evolutionary algorithm with host-parasite interpopulational relationship (Frank et al. 1993). In competitive co-evolution, advertisers' bidding strategies are optimized through arm races resulting from interactions and competitions among populations. Typically, co-evolutionary simulation consists of the following three major steps: evolutionary encoding, fitness evaluation, and genetic operations.

3.2.1. Evolutionary encoding

Encoding an SSA mechanism involves strategy encoding and mechanism encoding. We employ the real-coding scheme to encode the pure strategy space of each advertiser into a strategy population with a finite number of individual strategies. The initial strategy populations are generated by uniform sampling in advertisers' strategy spaces. An individual strategy is encoded as a chromosome with a finite number of gene locuses. The gene on each gene locus represents a component of an individual strategy, and all possible values for a gene are encoded as alleles. For instance, an individual strategy of the FSP auctions consists of two gene locuses, one containing the bid value and the other containing the choice to pay by first or second price scheme.

An SSA auction mechanism consists of the auction protocol and multiple dialogue rules. The former determines the sequence of advertisers' moves, and the latter the rules governing the auctions. These rules in most cases cover bid ranking and payment schemes. In co-evolutionary simulation, the auction protocol is encoded as the interaction protocol among strategy populations, while the dialogue rules are used to evaluate the fitness of individual strategies.

3.2.2. Fitness evaluation

The strategy populations of competing advertisers co-evolve with host-parasite relationships (Frank et al. 1993). Each strategy population takes turns to be a host and others are parasites in the fitness evaluation step. When evaluating the fitness of a host strategy, one individual strategy chosen from every other population will be selected as a parasite and matched with the host strategy to play an SSA auction game, as is shown in Fig. 1. The cumulative revenue earned by the host strategy in all possible host-parasite encounters can be considered as its fitness measure.

3.2.3. Genetic operations

Each advertiser's strategy population evolves through genetic operations. We use fitness proportionate selection, stochastic crossover, and mutation operators to evolve individual strategies. More specifically, selection is performed using the elitist roulette wheel scheme, which first migrates the fittest individuals in parent populations directly to offspring populations, and then performs the roulette wheel selections repeatedly until the offspring populations are generated. Moreover, each individual strategy stochastically crossovers with other strategies, and will be replaced by another strategy with a certain mutation rate.

3.3. Niche-based co-evolutionary mechanism design

Generally speaking, due to the genetic drift effect, the standard co-evolutionary algorithm will converge uniformly at one equilibrium (which possesses evolutionary stability and dynamic attainability properties) but miss all other alternative equilibria. Most SSA auction mechanisms, however, are associated with an infinite equilibrium continuum. To address this problem, we have developed a new co-evolutionary algorithm that incorporates the isolated niche technique. Using this approach, stable sub-populations are formed and maintained, and multiple equilibria evolve simultaneously.

Inspired by observed geographical isolation of species in nature, the isolated niche technique separates a population into several independently evolving niches, among which gene communications are not allowed (Rundle and Nosil 2005). As such, the effect of genetic drift is confined within each niche and global convergence will be avoided. Isolated niches can effectively maintain population diversity and increase the search efficiency for multiple solutions. This technique has been successfully used to solve the multi-modal function optimization and multi-issue bilateral negotiation problems (Li and Kang 2003, Yuan and Liang 2007).

The basic idea of the niche-based co-evolutionary SSA mechanism design is as follows. We search for all possible equilibrium

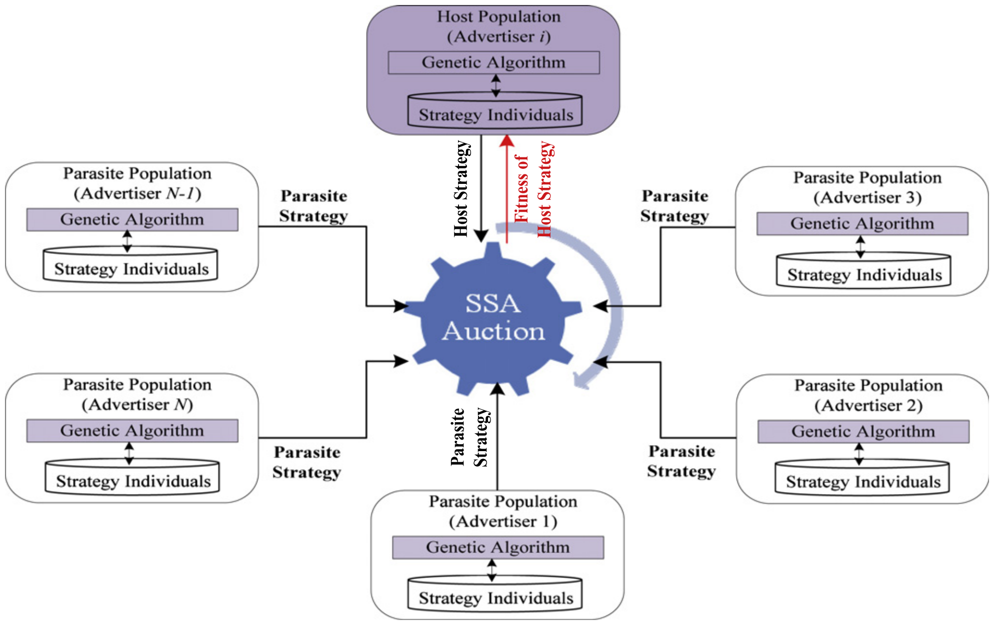

Fig. 1. Fitness evaluation in co-evolutionary simulation.

Table 1
Parameters setting the co-evolutionary simulation.

Auction parameters		Algorithm parameters	
Parameter	Value	Parameter	Value
Advertiser number	$N = 5(3)$	Maximum generation	$MaxG = 500$
Slot number	$K = 3(2)$	Population Size	$S = 20$
Per-click values	{0.9,0.7,0.5,0.3,0.1} ({0.9,0.6,0.3})	Niche number	$\gamma = 1(50)$
		Crossover rate	$crs = 0.6$
CTRs	{0.8,0.6,0.4} ({0.8,0.5})	Mutation rate	$mut = 0.02$
		Maximum encounters	$Maxtest = 20$
		Population variance	$\theta_{SPV} = 0.05$
		Mobile variance	$\theta_{EMV} = 0.001$
		Maximum niche score	$\theta_{max} = 20$
		Minimum niche score	$\theta_{min} = 1$

bidding behavior in advertisers' joint strategy space, which typically can be represented by an N-dimensional bid space R^N. For simulation simplicity, we normalize an advertiser's bid range to be within 0 and 1. The joint strategy space is then reduced to $[0,1]^N \in R^N$. In order to form niches in advertisers' populations, we define a search precision parameter $\gamma \geqslant 1$, and divide the strategy space of each advertiser into γ intervals. Under this partitioning scheme, advertisers' joint strategy space is divided into γ^N subspaces, each representing a niche in the co-evolving populations.[3] We generate N co-evolving sub-populations, each belonging to an advertiser's strategy population, to search for the optimal strategy profile within a niche. There are S individual strategies in each sub-population, and each individual strategy is a bid. Standard genetic operators are used to search the un-explored strategy space inside a niche, and no genetic communication (e.g., gene migration) is allowed among niches.

To determine evolutionary stability inside a niche within each generation, we propose two new measures, the static population variance (*SPV*) and the evolutionary mobile variance (*EMV*). More specifically, *SPV* is the variance of all individual strategies in a specific population, while *EMV* is the variance of the average bids during a number of generations in co-evolution. A strategy population will stably converge if these two indicators converge to zero. Computationally, in case when these two indicators drop below the preset thresholds, θ_{spv} and θ_{emv}, respectively, for all populations in a niche, we will consider the resulting populations as in an equilibrium state. The simulation bookkeeping module then saves advertisers' equilibrium bids and increments the score of this niche by one. The score of the niche is defined as the number of times an equilibrium is reached within this niche so far in the simulation runs.

The niches will be dynamically merged and divided to accelerate the co-evolution process. We define two scoring thresholds, θ_{max} and θ_{min}. If the score of a niche exceeds θ_{max}, we divide it into two sub-niches, and generate sub-populations for the new niches. In case when the score of a niche is below θ_{min}, we merge this niche with neighboring niches, and generate new sub-populations for the merged niche. These operations help increase the adaptability and speed of the co-evolution process.

The detailed co-evolutionary simulation algorithm following the ideas discussed above is as follows.

[3] Obviously, if we set $\gamma = 1$, the niche-based co-evolution falls back to the standard co-evolution, which searches advertisers' joint strategy space for a globally optimal equilibrium.

Algorithm 1. Niche-based co-evolutionary simulation for SSA mechanism design.

Input: An SSA auction mechanism.
Output: The equilibrium continuum of the input auction mechanism.
1 Parameters initialization;
2 Set Evolutionary generation $T = 0$;
3 For each advertiser $i \in [1,N]$ do//niche formation
4 Split the strategy space of advertiser i into γ intervals
5 Create a sub-population with S individual strategies for advertiser i in each interval
6 End for//resulting in γ^N niches in the joint strategy space
7 Repeat
8 For each niche, do
9 For each population $Pop_i \in Pop$ of advertiser i do// fitness evaluation
10 For each individual strategy $S_j \in Pop_i$ do//Host strategy
11 For $Test = 1$ to $MaxTest$ do// Generate $MaxTest$ encounters
12 For each population $Pop_k \in Pop$ such that $k \neq i$ do
13 Choose a strategy individual in Pop_k// parasite strategy
14 End for
15 Match the host and parasite strategies, and play SSA auctions
16 Calculate the revenue of the host strategy in the auction
17 End for
18 Fitness of S_j= discounted sum of revenues over $MaxTest$ encounters
19 End for
20 End for
21 End for
22 For each niche $n \in [1,\gamma^N]$, do
23 If $SPV_n < \theta_{SPV}$ and $EMV_n < \theta_{EMV}$ do
24 Save the average bids in all populations of the niche as an equilibrium
25 $scr_n = scr_n + 1$;//Increase the score of the niche by 1;
26 End if
27 If $T > MaxG/2$ do//merge and divide niches
28 If $scr_n < \theta_{min}$ do: merge the niche with neighboring niche.
29 If $scr_n > \theta_{max}$ do: divide the niche into two new niches.
30 End if
31 End for
32 For each population $Pop_i \in Pop$ do//genetic operations
33 Select individuals with elitist tournament scheme based on fitness
34 Apply crossovers and mutations with probabilities crs and mut
35 Produce offspring population
36 End for
37 Set evolutionary generation $T = T + 1$
38 Until ($T > MaxG$)
39 Output all the saved equilibria.

Note that the auction mechanism (as input) in Algorithm 1 is processed only in the SSA auction simulation step (Line 15). Typically, an SSA auction mechanism consists of an auction protocol and two dialogue rules, i.e., a ranking rule and a payment rule. The protocol and dialogue rules can be flexibly configured and easily integrated into our algorithm. For instance, switching from the GFP to GSP mechanism can be easily implemented by simply adjusting each advertiser's payment from her own bid to the bid of the next highest advertiser. In this sense, our algorithm is largely insensitive to the auction mechanisms under investigation.

3.4. Computational complexity analysis

In this section, we present a computational complexity analysis of the niche-based co-evolutionary simulation algorithm developed in the previous section. We decompose this algorithm into four parts to analyze its computational complexity. First, let's consider niche formation and initial population generation (Lines 3–6). In this part, we need to complete $\gamma * S$ random sampling operations to create a sub-population for each niche in an advertiser's population, and in total $N * \gamma * S$ operations for N advertisers. The time complexity is $O(N\gamma)$ since S is an constant in the algorithm. Second, we focus on fitness evaluation (Lines 8–21). There are altogether $MaxG * \gamma^N * N * S * Maxtest$ encounters for fitness evaluation during co-evolution. In each encounter, we select $N-1$ parasite strategies ($O(N)$), and simulate the SSA auction by ranking advertisers by their bids ($O(NlogN)$), and finally determine their revenue ($O(N)$). Therefore, the time complexity of fitness evaluation is $O(MaxG * \gamma^N * N * S * Maxtest * N * logN)$, or $O(\gamma^N N^2 logN)$. Third, consider niche operations (Lines 22–31). In each niche, calculating SPV and EMV, as well as merging or dividing niches, need $O(NS)$ time. Thus, the time complexity in this part is $O(MaxG * \gamma^N * NS)$, or $O(\gamma^N N)$, equivalently. Fourth, we concentrate on genetic operations (Lines 32–36). The time complexity for genetic operations is $O(MaxG * N * S)$ or $O(N)$, equivalently. To summarize, the overall computational complexity of our algorithm is $O(\gamma^N N^2 logN)$.

At the first glance, the computing time increases exponentially with the number of advertisers, which is to be expected of searching for equilibria in the high-dimensional joint strategy spaces of advertisers. However, in evaluating SSA mechanisms, this exponential complexity can be significantly reduced for the following reasons. First and most importantly, the performance of an SSA mechanism is an inherent property and do not vary with respect to the number of advertisers and niches. For instance, as predicted by theoretical models, the VCG mechanism is incentive compatible regardless how many advertisers are involved in auctions (Jansen and Mullen 2008). In this case, we can simply evaluate an SSA mechanism by running our algorithm for only one time with a randomly selected N and γ. As a result, the computational complexity will be reduced to $O(\gamma^N)$ with N as a constant. As such, the algorithm complexity depends only on the search precision of SSA equilibrium, controlled by γ. Second, as a mechanism evaluation tool, our algorithm does not necessarily need to run in an online and real-time environment. Meanwhile, the nature of simultaneous searches on niches makes parallel implementation possible through mechanisms such as map-reduce. Third, dynamic merging of non-equilibrium niches in our algorithm can significantly accelerate the co-evolution process.

4. Analyzing SSA mechanisms

In this section, we analyze four SSA auction mechanisms including GFP, GSP, VCG, and FSP, with co-evolutionary simulations. These mechanisms with the exception of FSP have been intensively

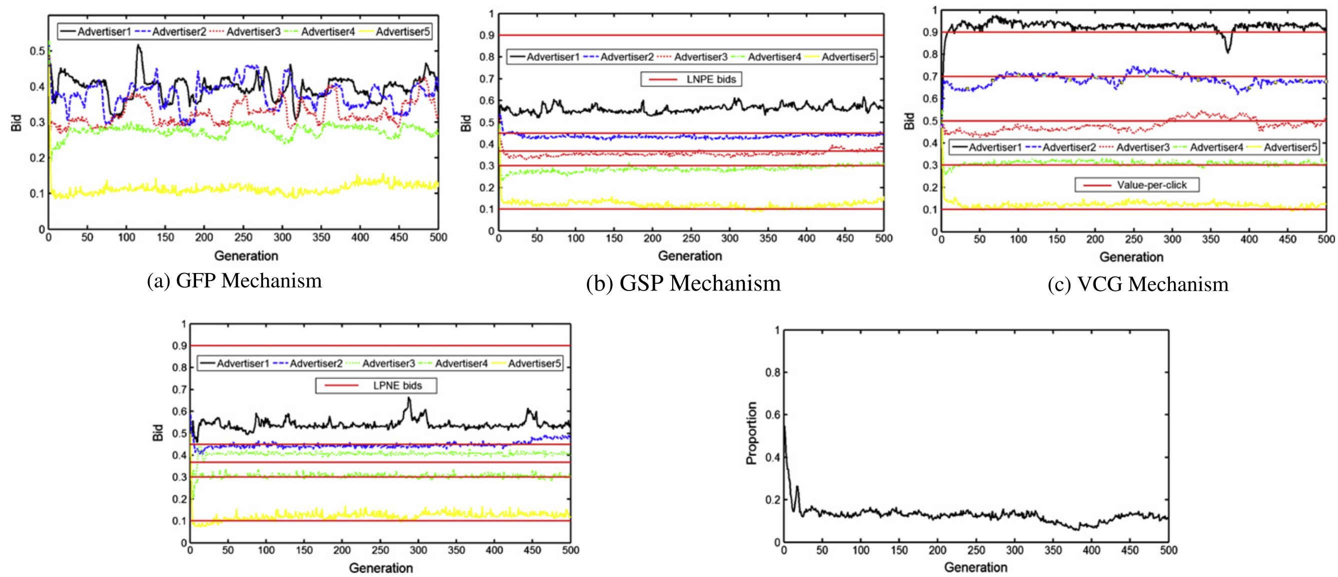

(a) GFP Mechanism (b) GSP Mechanism (c) VCG Mechanism (d) FSP Mechanism (e) Advertisers Preferring First-price Scheme

Fig. 2. Co-evolutionary dynamics of advertisers' bids.

studied in the literature. All mechanisms with the exception of VCG have been used in practice. Nevertheless, the equilibrium continuum of advertisers' bidding behavior and the performance of these mechanisms are not yet fully explored. Our research goal is to investigate the globally stable equilibria and the entire equilibrium continuum in SSA auctions with each of these four mechanisms, and then provide a qualitative and quantitative evaluation and comparison of the key properties of these mechanisms.

We proceed in two steps. First, we investigate the globally stable equilibria in each mechanism using the standard co-evolutionary algorithm. In standard co-evolution, advertisers' strategy populations uniformly converge at an optimal and stable equilibrium in advertisers' joint strategy spaces, and this equilibrium can be considered as the long-term outcome of SSA markets with revenue-maximizing advertisers. Second, we derive the equilibrium continuum using the niche-based co-evolutionary algorithm, and evaluate the performance of these auction mechanisms based on the derived equilibrium continuum. In niche-based co-evolution, local search will be performed within each niche, and result in all possible equilibria in SSA auctions.

4.1. Bidding behavior analysis based on standard co-evolutionary algorithm

We use the standard co-evolutionary algorithm (by setting $\gamma = 1$ in Algorithm 1) to search the global strategy space of each SSA mechanism for the evolutionarily stable equilibrium. Obviously, search engine companies prefer SSA mechanisms possessing evolutionary stability since auction outcomes can be easily and precisely predicted with advertisers' bids converging at a specific equilibrium.

The detailed setting of experimental parameters, including auction parameters and algorithm control parameters, are presented in Table 1. We ran co-evolutionary simulations for all four auction mechanisms. The detailed experimental results, including the evolutionary dynamics and convergence of advertisers' bids and fitness, are shown in Figs. 2–4.

We summarize the findings from Fig. 2 concerning bidding behavior.

(1) Advertisers' bids in GFP auctions do not converge. Furthermore, even their ranks change frequently during the co-evolution (Fig. 2a).
(2) The bids in GSP auctions stably converge to the LPNE (Fig. 2b), which indicates that LPNE is a globally optimal equilibrium in GSP auctions.[4] As such, SSA markets with a GSP auction mechanism will stabilize at an outcome in which all advertisers submit LPNE bids and obtain the maximum revenue in SNE.
(3) VCG auctions will result in a truthful-bidding equilibrium, since advertisers' bids converge to their private values-per-click (Fig. 2c).
(4) Interestingly, we observe that in FSP auctions, advertisers choosing to pay by the first-price scheme become (approximately) extinct during co-evolution (Fig. 2e). This indicates that the FSP mechanism will finally evolve to, and in fact is strategically equivalent to, the GSP mechanism. As a long-run stable outcome of FSP auctions, rational advertisers will learn to pay by second-price scheme and their bids will converge to the LPNE (Fig. 2d).

Fig. 3 indicates that GFP is not evolutionarily stable with the indicators *EMV* far exceeding the predefined thresholds. In contrast, other mechanisms possess evolutionary stability.

Finally, we observe from Fig. 4 that advertisers participating in GFP auctions cannot obtain stable revenue, while advertisers in GSP and FSP auctions will get the LPNE revenue and those in VCG auctions will get the truth-telling VCG revenue. Moreover, we observe from the simulation that the LPNE revenue equals the truth-telling VCG revenue.

4.2. Mechanism evaluation based on niched co-evolutionary algorithm

We now investigate the equilibrium continuum using the niche-based co-evolutionary algorithm. For illustration purposes, we consider an SSA auction scenario with three advertisers with values $\{0.9, 0.6, 0.3\}$ competing for two slots with CTRs $\{0.8, 0.5\}$. The niche parameter γ is set to 50. Other parameters are the same with those in Table 1. As can be seen from experimental results, the unstable GFP mechanism has an empty equilibrium continuum. Meanwhile, the strategically equivalent GSP and FSP mechanisms

[4] The reason why the top advertiser's bids do not converge precisely at the LPNE bid is that her bid values have no influence on all advertisers' revenues.

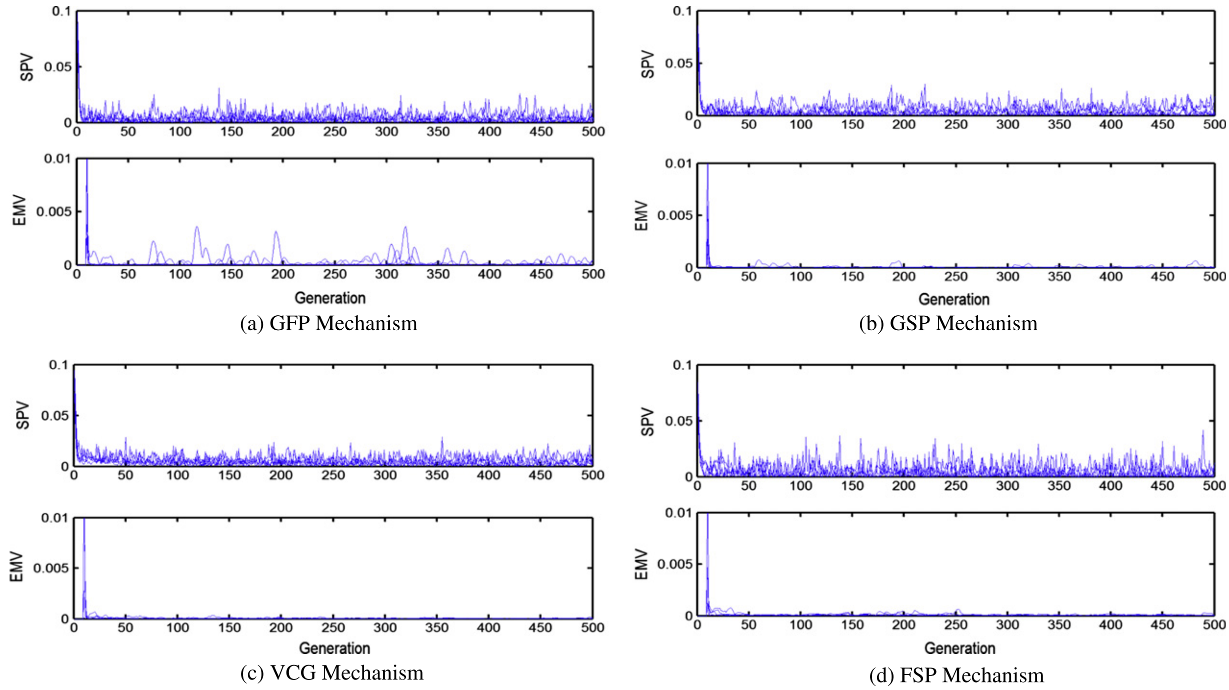

Fig. 3. Co-evolutionary stability of advertisers' bids.

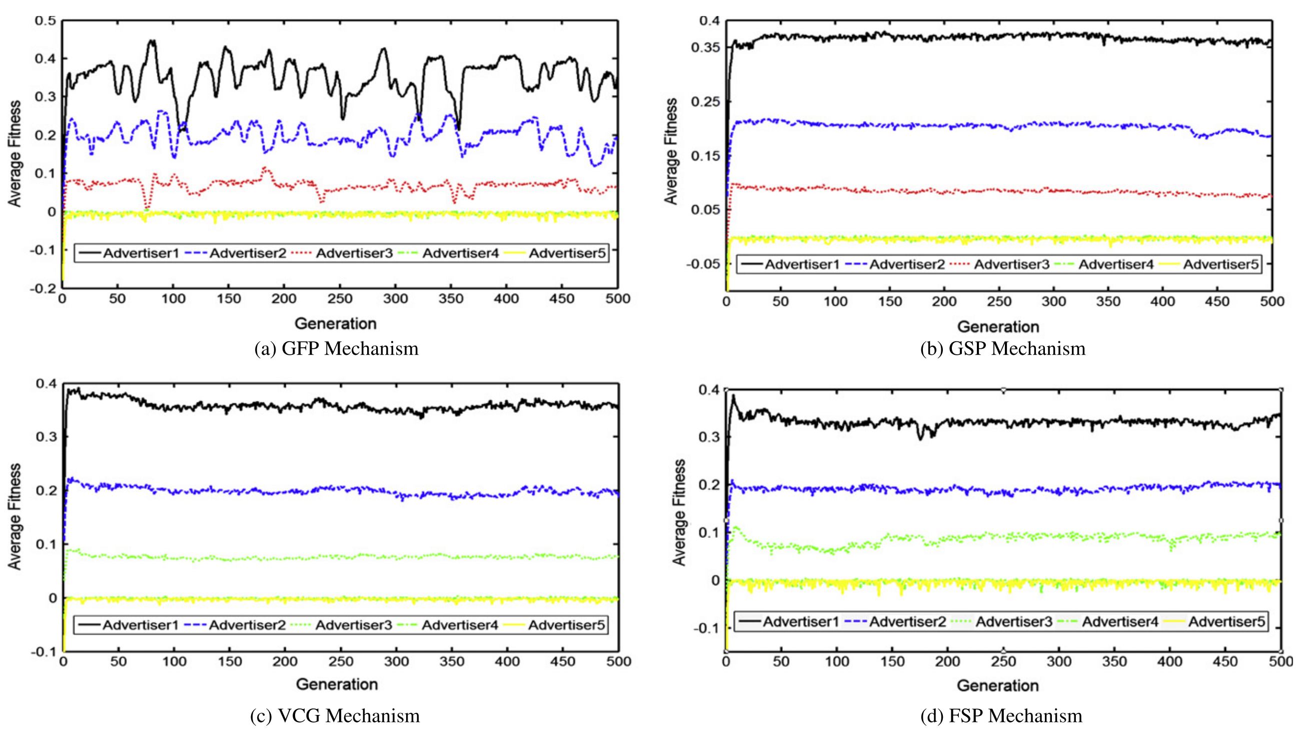

Fig. 4. Advertisers' fitness in co-evolution.

have similar equilibrium continuums. Therefore, we present in Fig. 5 only the equilibrium continuums of the GSP and VCG mechanisms.

It can be concluded from Fig. 5a that enabled by the niche technique, we have derived two clusters of equilibria in GSP auctions. One cluster, to the left, corresponds to non-output-truthful equilibria. The other, to the right, corresponds to output-truthful equilibria. Note that using the standard co-evolution, only a single optimal equilibrium in the output-truthful cluster can be derived. From Fig. 5b, we note that advertiser 3 will always be a loser in all equilibria. Advertiser 2 wins the top slot only in cases when the competing first advertiser submits a bid less than her private value (0.6). We can also approximately locate the upper bound and lower bound of the equilibrium continuum, namely, $\{0.9885, 0.7068, 0.5912\}$ and $\{0.3125, 0.3013, 0.0068\}$, which are the most preferred equilibria by the search engine company and advertisers, residing on the rightmost and leftmost sides of the 2-dimensional plane, respectively. These two extreme equilibrium

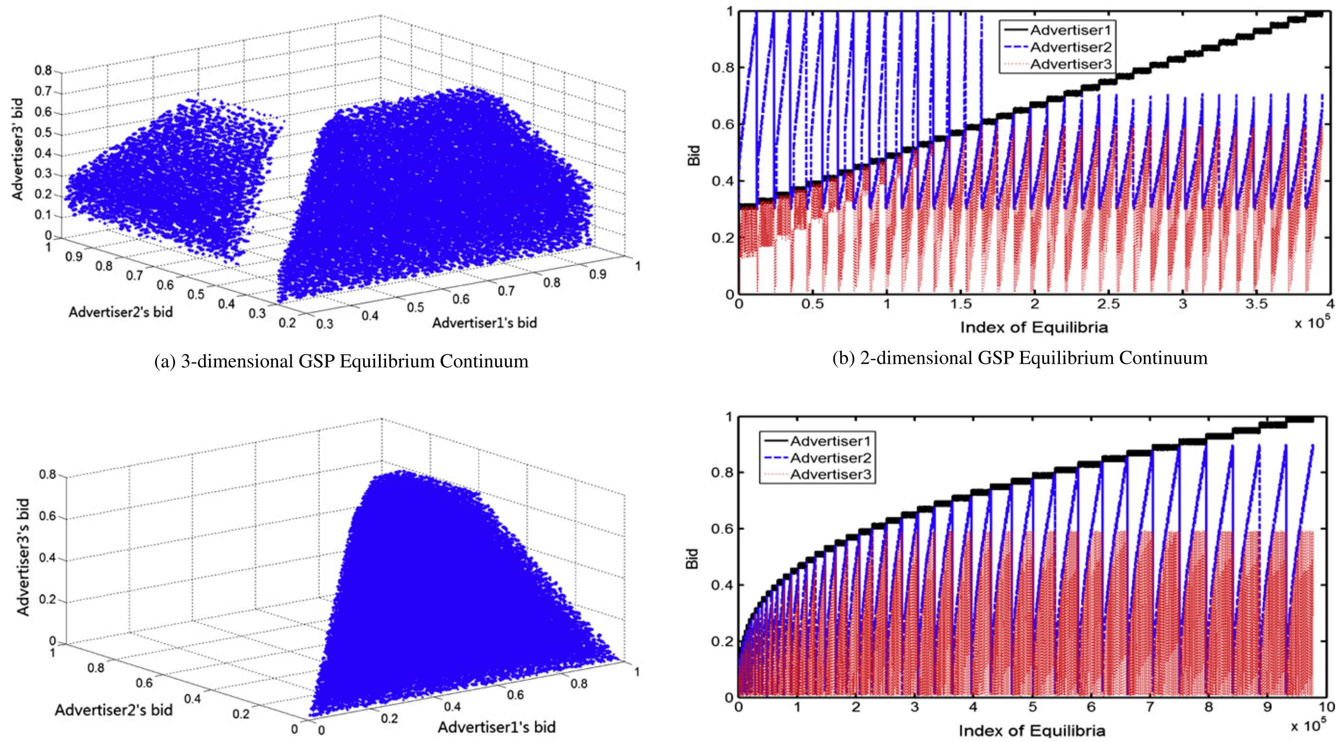

(a) 3-dimensional GSP Equilibrium Continuum

(b) 2-dimensional GSP Equilibrium Continuum

(c) 3-dimensional VCG Equilibrium Continuum

(d) 2-dimensional VCG Equilibrium Continuum

Fig. 5. Equilibrium continuums for GSP and VCG mechanisms.

bids help compute the performance measure of revenue ratio of the GSP mechanism. Analogously, we note from Fig. 5c and d that VCG is an output truthful auction mechanism with only one equilibrium continuum, in which advertisers are always ranked in decreasing order of their bids.

Based on these equilibrium continuums, we can now evaluate the key properties of SSA mechanisms. It is worth noting that since no equilibrium exists in GFP auctions, we assume that all bids satisfying individual rationality (i.e., not exceeding advertisers' private values) occur with equal possibility. We use all individually rational bids to evaluate the GFP auction mechanism.

The performance measures of the SSA mechanisms under investigation are listed in Table 2. We draw the following conclusions. First, the FSP mechanism performs equally well to the prevailing GSP mechanism. Considering the similarity in advertisers' bidding behavior, we can conclude that FSP is essentially equivalent to GSP. In other words, the hybrid FSP mechanism invented by Sogou.com, is in fact not a fundamentally new mechanism from a mechanism design perspective.

Second, VCG can maintain the maximum market efficiency in all equilibria (i.e., 1.02), while GFP might lead to the minimum market efficiency (i.e., 0.54). GSP and FSP can be considered as a tradeoff on average market efficiency between the optimal VCG and worst GFP mechanisms.

Third, all mechanisms have approximately equal minimum revenue ratio, while GFP and VCG may achieve the maximum revenue rate. On average, advertisers will obtain the largest share of auction surplus in VCG and a relatively high share in GFP auctions (although not stable). In contrast, search engines will prefer GSP and FSP auctions in which advertiser only obtains no more than half of auction surplus.

Fourth, all VCG equilibria are output truthful, while other mechanisms will lead to non-output truthful equilibrium bids.

Finally, as to qualitative measures, we can see that only VCG possesses incentive compatibility, and all mechanisms other than GFP are evolutionarily stable and may evolve to one equilibrium.

To summarize, VCG proves to be the best mechanism for both advertisers and the SSA markets among all the mechanisms studied. However, we experimentally observe that there is no revenue guarantee for the search engine in VCG auctions as advertisers may get all auction surpluses ($RR_{max} \rightarrow 1$) by strategic collusions. As a good alternative, search engines in GSP and FSP auctions will have

Table 2
Performance measures of SSA auction mechanisms.

Mechanism	Number of equilibria	Market efficiency			Revenue ratio (%)			Output truthfulness (%)	Incentive compatibility	Evolutionary stability
		ME_{max}	ME_{min}	ME_{avg}	RR-max	RR_{mln}	$RRavg$			
GFP	0	1.02	0.54	0.9074	99.98	15.48	62.54	44.42	No	No
GSP	394,721	1.02	0.93	1.0012	76.04	15.24	49.10	79.16	No	Yes
VCG	977,598	1.02	1.02	1.02	99.22	15.46	64.42	100	Yes	Yes
FSP	394,835	1.02	0.93	1.0011	76.14	15.5	49.14	79.03	No	Yes

guaranteed revenues in case when advertisers submit the LPNE bids. We strongly believe that this might be one of the factors behind the huge practical success of GSP.

4.3. Discussion of experimental results

The experimental results readily support the existing theoretical analyses. For instance, it has been proved by theoretical and empirical analysis that advertisers bidding in GFP auctions will not form a equilibrium (Zhang and Feng 2005, Edelman and Ostrovsky 2007). This prediction has been validated by the co-evolutionary simulation results shown in Fig. 2a. Edelman et al. (2007) theoretically proved that in GSP auctions, advertisers' bids will stabilize at a ''locally envy free'' equilibrium, which is equivalent to the LPNE. All players, including search engines and advertisers, will get the same payoff in the LPNE as in auctions with the VCG mechanism. This has been observed in Fig. 2b that advertisers' bids stably converge to the LPNE values. Meanwhile, VCG has been proved to be a incentive compatible mechanism with a truth-telling equilibrium in dominant strategies (Aggarwal et al. 2006). This can be verified in Fig. 2c that advertisers' bids in the VCG mechanism have stably converged to their private values. We can conclude that co-evolutionary simulation is effective in analyzing advertisers' equilibrium bidding behavior in SSA auctions.

Our simulation also offers novel insights about SSA mechanism design and evaluation. For instance, we have validated that the FSP mechanism is strategically equivalent to the prevailing GSP mechanism. Through discovering all possible equilibria in SSA auctions, we have determined the distribution of the equilibrium continuum in the GSP and VCG mechanism, and investigated the key performance measures of these auction mechanisms. As a key result, we experimentally observe that advertisers can squeeze out all surplus in SSA auctions with GFP and VCG mechanisms through strategic bidding, which has not been explored by the existing theoretical research and can be considered as one of the major reasons why search engines do not use these mechanisms. These simulation results can effectively complement and enhance the existing analytical results, and offer new insights on SSA auctions with more complex mechanisms.

5. Conclusion and future work

This paper focuses on SSA mechanism design and evaluation. To address the limitations of the analytical mechanism design framework, we propose a niche-based co-evolutionary simulation approach for SSA mechanism design, aiming at computationally analyzing advertisers' equilibrium bidding behavior, and deriving the key performance measures of various kinds of SSA auction mechanisms.

Our approach has the following managerial implications. For online advertisers, our work can help generate the optimal equilibrium bids. For Web search engine companies, it can help better understand advertisers' bidding behavior and dynamics through analyzing the equilibrium continuum of SSA auctions. It can also be used to evaluate key performance of alternative SSA auction mechanisms.

Our research has two major limitations. First, the proposed approach can only be used to evolve advertisers' equilibrium bidding behavior in SSA scenarios where advertisers have complete information, or reliable estimations about all advertisers' per-click values and CTRs for all slots. Second, co-evolutionary simulations for SSA auctions with a large number of advertisers will impose major computational overhead.

In the future work, we plan to extend our approach to handle incomplete information settings through maintaining a hybrid strategy population for each advertiser according to her Bayesian beliefs of the competitors' private per-click values. To deal with the computational overhead, we are working on parallelizing the key simulation algorithm. We also plan to conduct formal analysis of the co-evolutionary simulation framework through replicator dynamics theory.

Acknowledgments

We gratefully acknowledge funding support from the National Natural Science Foundation of China (Grant 71102117, 70890084, 71025001, 71071152, 60921061) and the Chinese Academy of Sciences (Grant 2F07C01). We would like to thank the Editor-in-Chief, the special issue Editors, and anonymous reviewers for their constructive comments and suggestions in improving this study.

References

Aggarwal, G., Goel, A., and Motwani, R. Truthful auctions for pricing search keywords. In Proceedings of the 7th ACM Conference on Electronic Commerce, June 11–15, Ann Arbor, Michigan, 2006, 1–7.

Aggarwal, G., Muthukrishnan, S., Pal, D., and Pal, M. General auction mechanism for search advertising. In Proceedings of the 18th International Conference on World Wide Web, April 20–24, Madrid, Spain, 2009, 241–250.

Apaloo, J., Muir, P. W., and Hearne, J. W. Multi-species evolutionary dynamics. *Evolutionary Ecology*, 19, 2005, 55–77.

Balcan, M. F., Blum, A., Hartline, J. D., and Mansour, Y. Sponsored search auction design via machine learning. Technical report CMU-CS-050143, 2005.

Bu, T. M., Li, L., and Qi, Q. On robustness of forward-looking in sponsored search auction. *Algorithmica*, 58, 2010, 970–989.

Ciaramita, M., Murdock, V., and Plachouras, V. Online learning from click data for sponsored search. In Proceedings of the 17th International Conference on World Wide Web, April 21–25, Beijing, China, 2008, 227–236.

Ceppi, S., Gatti, N., and Gerding, E. Mechanism design for federated sponsored search auctions. In Proceedings of the 25th Conference on Artificial Intelligence, August 7–11, San Francisco, California, 2011.

Cliff, D. Explorations in evolutionary design of online auction market mechanisms. *Electronic Commerce Research and Applications*, 2, 2003, 162–175.

Dash, R. K., Parkes, D. C., and Jennings, N. R. Computational mechanism design: a call to arms. *IEEE Intelligent Systems*, 18, 6, 2003, 40–47.

Deng, X. T., Sun, Y., Yin, M., and Zhou, Y. H. Mechanism design for multi-slot ads auction in sponsored search markets. *Lecture Notes in Computer Science*, 6213, 2010.

Dutting, P., Henzinger, M., and Weber, I. An expressive mechanism for auctions on the web, In Proceedings of the 20th International Conference on World Wide Web, March 28–April 1, Hyderabad, India, 2011, 127–136.

Edelman, B., and Ostrovsky, M. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43, 2007, 192–198.

Edelman, B., Ostrovsky, M., and Schwarz, M. Internet advertising and the generalized second price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97, 1, 2007, 242–259.

Feng, J., Bhargava, H. K., and Pennock, D. M. Implementing sponsored search in web search engine: computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 19, 1, 2007, 137–148.

Frank, S. A. Evolution of host-parasite diversity. *Evolution*, 47, 6, 1993, 1721–1732.

Garg, D., and Narahari, Y. An optimal mechanism for sponsored search auctions on the web and comparison with other mechanisms. *IEEE Transactions on Automation Science and Engineering*, 6, 4, 2009, 641–657.

Iyengar, G., Phillips, D., and Stein, C. Models of malicious behavior in sponsored search. In Proceedings of the 2007 Spring Simulation Multiconference, March 25–29, Norfolk, Virginia, USA, 2007, 143–151.

Jansen, B. J., and Mullen, T. Sponsored search: an overview of the concept, history and technology. *International Journal of Electronic Business*, 6, 2, 2008, 114–131.

Li, Y., and Kang, Z. A. Two-level subspace evolutionary algorithm for solving multi-modal function optimization problems. *Wuhan University Journal of Natural Sciences*, 8, 1, 2003, 249–252.

Munsey, M., Veilleux, J., Bikkani, S., Teredesai, A., and Cock, M. D. Born to trade: a genetically evolved keyword bidder for sponsored search. In Proceedings of the IEEE World Congress on Evolutionary Computation, July 18–23, Barcelona, Spain, 2010.

Pardoe, D., Stone, P., Tsechansky, M. S., Keskin, T., and Tomak, K. Adaptive auction mechanism design and the incorporation of prior knowledge. *INFORMS Journal on Computing*, 22, 3, 2010, 353–370.

Phelps, S., McBurney, P., Parsons, S., and Sklar, E. Co-evolutionary auction mechanism design: a preliminary report. *Lecture Notes in Computer Science*, 2531, 2002, 123–142.

Rundle, H. D., and Nosil, P. Ecological speciation. *Ecology Letters*, 8, 2005, 336–352.

Vijay, K., and Perry, M. Efficient mechanism design. In *Mimeo*, Penn State University and Hebrew University, 1997.

Weibull, J. W. *Evolutionary Game Theory*. MIT Press, Cambridge MA, London, 1995.

Yuan, Y., and Liang, Y. Q. Multi-issue negotiation research based on niched co-evolutionary genetic algorithm. In Proceedings of the Eighth ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing, July 30–August 1, Qingdao, China, 2007, 564–569.

Yuan, Y., Zeng, D., Li, L. J., and Bao, N. Bidding behavior and equilibrium on competitive keyword advertising markets. In Proceedings of the 21st Workshop on Information Technologies and Systems, December 3–4, Shanghai, China, 2011.

Zhang, X. M., and Feng, J. Price cycles in online advertising auctions. In Proceedings of the 26th International Conference on Information Systems, December 11–14, Las Vegas, Nevada, USA, 2005.

Zhou, Y., and Lukose, R. Vindictive bidding in keyword auctions. In *Proceedings of the Second Workshop on Sponsored Search Auctions, in conjunction with the ACM Conference on Electronic Commerce (EC 2006), Ann Arbor, Michigan*, ACM Press, New York, 2006.