A Bidding Agent for Advertisement Auctions: An Overview of the CrocodileAgent 2010

Irena Siranovic, Tomislav Cavka, Ana Petric, and Vedran Podobnik

University of Zagreb, Faculty of Electrical Engineering and Computing
Zagreb, Croatia
{irenasiranovic,tomislavcavka1}@gmail.com
{ana.petric,vedran.podobnik}@fer.hr

Abstract. Sponsored search is a popular form of targeted online advertising and the most profitable online advertising revenue format. Online publishers use different formats of unit price auctions to sell advertising slots. In the Trading Agent Competition Ad Auctions (TAC/AA) game, intelligent software agents represent a publisher which conduct keyword auctions and advertisers which participate in those auctions. The publisher is designed by game creators while advertisers are designed by game entrants. Advertisers bid for the placement of their ads on the publisher's web page and the main challenge placed before them is how to determine the right amount they should bid for a certain keyword. In this paper, we present the CrocodileAgent, our entry in the 2010 TAC AA Tournament. The agent's architecture is presented and a series of controlled experiments are discussed.

Key words: trading agents, sponsored search, keyword auctions

1 Introduction

Internet advertising provides a significant income stream for online publishers. Revenue of major search engines such as Google[1], Yahoo![2] and MSN[3] amounts to tens of billions of dollars annually (e.g., in 2010, Google reported total advertising revenues over USD $28 billion[4]). According to the report[5] published by the Interactive Advertising Bureau[6] and PricewaterhouseCoopers LLP[7], sponsored search is the most profitable online advertising revenue format which accounted for 47% of the total Internet advertisement revenue in the USA in the first half of 2010. In sponsored search (i.e., keyword advertising) publishers (i.e., search engines) use different formats [1] of unit price auctions (e.g., keyword auctions) to sell advertising slots (i.e., positions in the list that contains search results) [2].

---

1 http://www.google.com/
2 http://www.yahoo.com/
3 http://www.msn.com/
4 http://investor.google.com/financial/tables.html
5 http://www.iab.net/media/file/IAB_report_1H_2010_Final.pdf
6 http://www.iab.net/
7 http://www.pwc.com/

Siranovic, Irena; Cavka, Tomislav; Petric, Ana; Podobnik, Vedran.
A Bidding Agent for Advertisement Auctions: An Overview of the CrocodileAgent 2010.
Lecture Notes in Business Information Processing. 119 (2013); 71-86.

In a keyword auction advertisers bid for the placement of their ads (i.e., rank of the ad in the results of the sponsored search) which are then displayed on the publisher's web page. The format of sponsored search results is very much alike the format of generic search results. Usually, it is comprised of the title of the ad, a short description and a hyperlink to the advertiser's web page or the web page of the advertised product. An ad is chosen for a specific keyword(s) from the user's query, thus targeting users interested in advertiser's products.

A publisher conducts keyword auctions and solicits bids. When a user submits a query containing one or more keywords, sponsored ads are displayed to the user alongside the results of a generic search mechanism. At the end of an auction, the publisher ranks advertisers' bids (i.e., determines the placement of their ads and the cost-per-click (CPC) for those ads) [3]. CPC is the price that an advertiser pays to the publisher each time its ad is clicked on.

According to earlier studies on user behaviour (e.g., [4]) the higher the position of the search result (e.g., ad, document) is, the users will be more likely to click on it. This phenomenon, where the probability that the result will be clicked depends not only on its relevance, but also on its position in the search results, is known as the position bias. Several models of the position bias have been proposed and the cascade model gave the best explanation for the position bias in early ranks [5]. However, the latest research has shown that 46% of the users do not click sequentially (i.e., start from the best ranked result and continue to lower ranked ones) and 57% of them do not behave as suggested by the cascade model (i.e., first click on the higher and afterwards on the lower positioned results) [6].

A challenge placed upon publishers evolves around the selection of a mechanism which will result with highest profits. There are two frequently used mechanisms for ranking solicited bids: i) rank-by-bid, and, ii) rank-by-revenue. As the name of the rank-by-bid mechanism states the bids are sorted in a descending order (i.e., the bids offering a higher CPC get higher ranked positions), while the rank-by-revenue mechanism multiplies the offered CPC with the ad's expected relevance (i.e., the percentage of users that will click on the ad once it is displayed to them) and afterwards sorts bids in a descending order of the calculated product [7]. In addition, some mechanisms offer advertisers the possibility to target a certain group of users by specifying the context for viewing ads (e.g., user's location, time of day) and to control exposure by limiting the number of times an ad should be displayed to users [8].

From the advertisers' point of view, the question is how to determine the right amount they should bid for a certain keyword(s) since the probability that their ad will be better ranked than other ads rises as they place higher bids [9]. The complexity of the answer to this question increases as the availability of information about user behaviour, as well as bidding behaviour of other advertisers decreases.

The paper is organized as follows. Section 2 describes the characteristics of the Trading Agent Competition Ad Auctions (TAC/AA) game. A brief overview of the research in the area of TAC/AA bidding strategies is given in Section 3. Section 4 presents the CrocodileAgent 2010, our entrant in the 2010 TAC/AA Tournament. Section 5 presents the conducted experiments and the obtained results, while Section 6 concludes the paper and gives an outline for future work.

2 TAC/AA game

Researchers test advertisers' bidding strategies by using the designed market simulators which provide a risk free environment [10]. The TAC/AA game [11, 12], which was released in 2009, is based on such a market simulator. In other TAC games market simulators are used to find perspective solutions for the supply chain management problem [13, 14, 15], market design problem [6, 16, 17] and energy trading in smart grid environments [18].

The TAC/AA game enables advertisers to bid on multiple queries and to define budget constraints in an information-lacking and competitive environment. Eight intelligent program agents which represent online advertisers participate in the game. Each advertiser sells 9 different products which are specified by the manufacturer (i.e., Flat, Lioneer and PG) and the component (i.e., TV, Audio and DVD). Furthermore, each advertiser is specialized for one manufacturer and for one component type.

The ad ranking mechanism varies between rank-by-bid and rank-by-revenue [19] and is chosen at the beginning of each game. Users can generate 9 different queries by specifying both manufacturer and component (i.e., F2 level query), 6 different queries by specifying only manufacturer or component (i.e., F1 level query) and one query where neither manufacturer nor component are specified (i.e., F0 level query). Correspondingly, the conversion probability increases as more keywords are specified in the query. Users can find themselves in one of the following states: i) non-searching, ii) searching, and, iii) transacted; while the user population in each state (and sub-state) is modelled as a Markov chain.

In the TAC/AA game scenario, when a user submits a query, an auction for the given keyword(s) starts. The auction is an instance of the repeated generalized second price auction (i.e., the price that the advertiser pays for the position of its ad is determined by the price which was offered by the winner of the next-best position in its bid) and the first position is allocated to the bidder with the highest bid. An advertiser sends a bid bundle (i.e., one bid for each query class) to the publisher every day. A bid consists of: i) the CPC an advertiser is ready to pay, ii) the chosen ad which can be generic or targeted (i.e., specified manufacturer and/or component), iii) budget limits for each query class, and, iv) budget limit for all queries altogether for the following day (optional). The publisher uses the information from the bids when it runs ad auctions for the received user queries. Advertisers receive daily reports about the outcomes of the prior (i.e., two days old) auctions and use those information to generate new bids. Daily reports include: i) query report, ii) account status report, and, iii) sales report. The game lasts 60 virtual days.

3 Related Work

The bidding strategy for keyword auctions has been a great challenge for researchers that conducted various simulations and empirical analyses on this matter. Berg *et.al.* [20] have presented autonomous bidding strategies for ad auctions which are based the on click probability and the CPC estimations. They use two kind of algorithms for bid optimization: i) rule-based algorithms, and, ii) greedy multiple choice Knapsack algorithms. Pardoe and Stone [21] have shown a particle filter that can be used for estimating other agents' bids given a periodic ranking of their bids. The particle filter is used for estimating bids of other advertisers. Since the information revealed about competing advertisers is limited, they have shown how such models can learn from the past bidding data. Cigler [22] has described several bidding strategies (Return-on-investment (ROI), Knapsack ROI, Balanced Best-response, online Knapsack) and evaluated their performance empirically. The best performing strategies were ROI and Balanced Best-response. Furthermore, Cigler presented a new profit maximizing strategy for multiple keyword ad campaigns that takes into account the budged constraint and has shown to be successful particularly for small budgets.

4 CrocodileAgent 2010

With an intention to better investigate strategic approaches to ad auctions mechanisms, an advertiser agent CrocodileAgent has been designed. A bidding strategy was chosen among several strategies (i.e., profit maximization, linear regression) [23] and analysed through controlled experiments. The CrocodileAgent's bidding model used for participating in ad auctions is shown in Figure 1 and can be divided into three logical segments: i) the ad generator, ii) the CPC generator, and, iii) the daily spend limit generator.

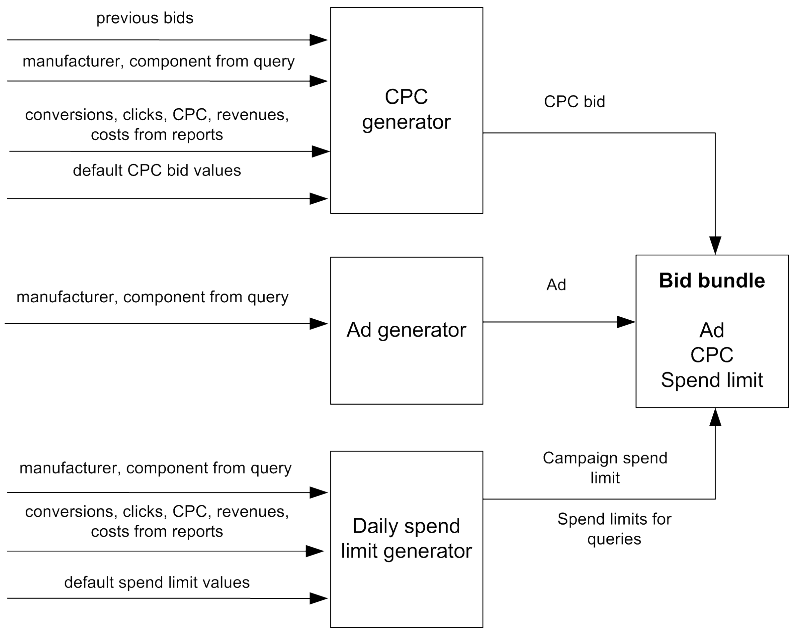

Fig. 1. CrocodileAgent's bidding model

4.1 The ad generator algorithm

The method chosen for ad generation is based on query focus levels and is shown in Algorithm 1. This method was chosen by analysing final results in game simulations with several different implemented algorithms for ad generation. The results shown that the conversion probability is the highest when a user submits a query with the focus level F2. Since there are 9 different products and each user has a preference for just one of them, it is highly probable that a user will click on the ad that matches the query when he/she submits an F2 query. Furthermore, we have concluded that if an agent provides the generic ad for the submitted F2 query, the probability that the user will click on the ad decreases significantly, especially if competing agents chose one of targeted ads.

Algorithm 1: CrocodileAgent's ad generation algorithm

```
if query focus level = F0 or F1 then
    generated ad = generic
else
    generated ad = targeted (F2 manufacturer, F2 component)
end
```

For a submitted F1 query, the CrocodileAgent generates generic ads – if the CrocodileAgent generated a targeted ad, it would have to decide which component or manufacturer to add alongside the specified keyword (i.e., manufacturer or component, respectively). The negative aspect of the latter approach is that the targeted ad has to match user preferences in order to accomplish a conversion (the probability that the generated ad matches user preferences is 1/3). If the ad does not match user preferences, the click probability decreases according to the targeting factor [11], and vice versa. Moreover, if the ad does not match user preferences, the user may click on the ad, but he/she will not convert and thus he/she will only increases the agent's cost. On the other hand, the benefit of the approach that results with a targeted ad manifests in case when the ad matches the submitted query – positive aspects of this approach are: i) additional profit gained by component through component specialization bonus (CSB) or manufacturer specialization bonus (MSB), ii) increased odds of converting based on CSB, and, iii) increased targeting effect.

4.2 The CPC generator algorithm

The CPC bid is defined for every query type with an intention to maximize the profit (profit = revenue - cost). Taking into consideration the two days old information in the received reports, the algorithm for generating the CPC for the first two days of the game contains fixed values of CPC. Values for all query types are defined based on matches between the component and manufacturer in a query and agent's specialties as shown in Algorithm 2.

Algorithm 2: CrocodileAgent's initial bidding algorithm

$result$ = match between query and agents specialties;
switch *result* **do**
  case *miss*
    $bid = \alpha;$
  case *miss-neutral*
    $bid = \beta;$
  case *miss-hit*
    $bid = \gamma;$
  case *neutral*
    $bid = \delta;$
  case *neutral-hit*
    $bid = \epsilon;$
  case *hit*
    $bid = \xi;$
endsw
where $\beta < \gamma < \alpha \leq \delta \leq \epsilon < \xi;$

The more accurate the match is, the higher the bid is. Additionally, the values are scaled by agent's capacity (i.e., the bid is slightly decreased in case of medium or low capacity). Definitions of matching values are listed in Table 1, while Table 2 contains the parameter values. All parameter values were determined by using a heuristic approach. The decisions were based on conducted game simulations (locally and in the TAC/AA 2010 qualifying rounds) by comparing the CrocodileAgent's profit for different sets of parameters used.

Table 1. Definitions of specialty matching values

Value	Query type	Definition
miss	F2	Component and manufacturer from query do not match agent's specialties.
miss-neutral	F1	Component or manufacturer from query does not match agent's specialties, second keyword is null.
miss-hit	F2	Component or manufacturer from query does not match agent's specialties, second keyword is a match.
neutral	F0	A query without specified manufacturer and component.
neutral-hit	F1	Component or manufacturer from query matches agent's specialties, second keyword is null.
hit	F2	Component and manufacturer from query are agent's specialties.

A method for defining the CPC bids in the remainder of the game is based on the calculations of the conversion rate (i.e., the ratio of average number of clicks to conversions). If the rate is satisfactory, the new bid is based on the CPC bid from two days

Table 2. Values of parameters for initial bidding

Parameter	α	β	γ	δ	ϵ	ξ
Value	1.15	0.65	1.05	1.15	1.15	1.25

ago. Otherwise, the new bid is based on the CPC that the agent actually paid. The bid is later adjusted depending on the query focus level and the CrocodileAgent's specialties. The CPC generator algorithm is shown in Algorithm 3.

Algorithm 3: CrocodileAgent's CPC generator algorithm

$revenue$ = last revenue for a query received in report;
n_{click} = average number of clicks per day;
$n_{conversion}$ = average number of conversions per day;
$query_{fl}$ = query focus level;
$result$ = match between query and agent's specialties;

if $revenue \neq 0$ **then**
  $mod = n_{click}/n_{conversion}$;
  bid = DefineBid(mod);
  if $query_{fl} == F0$ *or* $query_{fl} == F1$ **then**
    $bid = focus\ level\ parameter \cdot bid$
  end
  if $result == hit$ **then**
    $bid = specialty\ parameter \cdot last\ paid\ bid$
  end
else
  if $query_{fl} == F2$ **then**
    bid = min($minimum\ bid\ zero,\ hit\ parameter \cdot last\ bid_{day-2}$);
    if $manufacturer\ and\ component\ of\ query == agent's\ specialties$ **then**
      $bid = specialty\ parameter\ zero \cdot bid$
    end
  else
    $bid = parameter\ zero \cdot last\ bid_{day-2}$;
  end
end

function DefineBid(mod)
if $mod < ratio\ lower\ bound$ **then**
  bid = max($minimum\ bid,\ decrease\ factor \cdot last\ bid_{day-2}$);
else if $ratio\ lower\ bound \leq mod \leq ratio\ middle\ value$ **then**
  bid = max($minimum\ bid,\ increase\ factor \cdot last\ paid\ bid$);
else
  bid = max($minimum\ bid,\ max\ increase\ factor \cdot last\ paid\ bid$);
end
return bid;

When determining CPC bids the most important aspect is the ”quality” of an ad, which is measurable through the agent’s profit. The quality of an ad is defined as an estimation that the user’s click on that ad will turn into a conversion. The CrocodileAgent’s bidding strategy does not change during the game. Table 3 contains the parameters that gave the most satisfactory results in the conducted experiments.

Table 3. Values of bidding parameters during a game

Parameter	Definition	Value
ratio lower bound		10
ratio middle value		0.65
minimum bid	minimum CPC that an agent will bid	0.10
decrease factor	decrease factor in case of very good mod	0.97
increase factor	increase factor in case of average mod	1.15
max increase factor	increase factor in case of very bad mod	1.20
focus level parameter		0.80
speciality parameter		1.30
minimum bid zero	minimum CPC that an agent will bid in case that last revenue was zero	0.30
hit parameter		1.30
speciality parameter zero		1.20
parameter zero	a parameter used when agent’s last revenue was zero	1.05

4.3 The spend limit manager

General spend limit Too many conversions lead to decrease of the possible conversions in the future due to stock shortage, as defined by the game rules [12]. The stock management policy is necessary since it can happen that a certain amount of products must be immediately available in order to avoid significant profit loss. On the other hand, excessive product storage becomes ”dead capital”. Additionally, the other driver for using spend limits are the clicks generated by informational searchers.

The general spend limit manager algorithm adjusts the bid according to the agent’s capacity and it is shown in Algorithm 4, while corresponding parameter values are listed in Table 4. In the first five days of the game the limit is fixed and determined based on the agent’s capacity. After analysing the results of the controlled experiment we have concluded that the lack of spend limit has the most significant (negative) impact on CrocodileAgent’s profit when its capacity is low.

Query spend limit The query spend limit manager defines a distribution of agent’s investments in order to maximize its profit and ensures that the maximum number of clicks for a certain ad is limited in accordance with its predicted quality. The spend limits for the first two days are fixed and they are calculated based on the matching of the CrocodileAgent’s specialties with the keywords in a query. In the remainder

Algorithm 4: General spend limit manager

if *first five days* **then**
  switch *capacity* **do**
    case *low*
      $spend\ limit = spend\ limit\ low\ fixed$;
    case *medium*
      $spend\ limit = spend\ limit\ medium\ fixed$;
    case *high*
      $spend\ limit = spend\ limit\ high\ fixed$;
  endsw
else
  switch *capacity* **do**
    case *low*
      $spend\ limit = spend\ limit\ low$;
    case *medium*
      $spend\ limit = min(spend\ limit\ medium, custom\ profit)$;
    case *high*
      $spend\ limit = min(spend\ limit\ high, custom\ profit)$;
  endsw
end

Table 4. Values of parameters for general spend limit

Parameter	Min	Parameter	Min
spend limit low fixed	700	spend limit medium fixed	1000
spend limit high fixed	1200	spend limit low	750
spend limit medium	1150	spend limit high	1350

of the game the spend limit is defined according to the click and conversion ratio, as well as earlier profits. The low ratio corresponds with the successful advertisement. Therefore, the lower the ratio is, the higher the limit will be. The query spend limit manager algorithm is shown in Algorithm 5, while the corresponding parameter values are listed in Table 5.

Table 5. Values of parameters for general spend limit

Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
a	60	b	80	ratio lower bound	5	neutral	1.10
c	100	d	120	ratio middle value	10	neutral-hit	1.20
e	140	f	180	minimum limit	20	hit	1.30
				default limit	40		

Algorithm 5: Query spend limit manager

n_{click} = average number of clicks per day;
$n_{conversion}$ = average number of conversions per day;
mod = n_{click}/$n_{conversion}$;
$result$ = a match between query and agent's specialties;
if *first two days* **then**
    switch *result* **do**
        case *miss*
            $limit = a$;
        case *miss-neutral*
            $limit = b$;
        case *miss-hit*
            $limit = c$;
        case *neutral*
            $limit = d$;
        case *neutral-hit*
            $limit = e$;
        case *hit*
            $limit = f$;
    endsw
else
    if $mod < ratio\ lower\ bound$ **then**
        limit = min (minimum limit, revenue/2);
    else if $ratio\ lower\ bound \leq mod \leq ratio\ middle\ value$ **then**
        limit = min (minimum limit, revenue/3);
    else
        limit = default limi;
    end
    limit = `SpecialtyMatchingLimit` (*result, limit*)
end

function `SpecialtyMatchingLimit` (*result, limit*)

switch *result* **do**
    case *neutral*
        $limit = neutral \cdot limit$;
    case *neutral-hit*
        $limit = neutral - hit \cdot limit$;
    case *hit*
        $hit : limit = hit \cdot limit$;
endsw

5 Controlled experiment

In order to evaluate the performance of the CrocodileAgent, which placed 6th in the TAC/AA 2010 Competition Finals, an experiment was conducted by repeating games taking into consideration a fer distribution of agent capacity. Based on the analysis of the results, the CrocodileAgent's deficiencies were identified and guidelines for future improvements were set.

The participants in the experiment were the following agents which competed in the TAC/AA 2010 Competition: *TacTex*, *Mertacor*, *Schlemazl*, *CrocodileAgent*, *tauagent*, *EPFLAgent*. Additionally, due to the lack of TAC/AA 2010 agents in the official agent repository[8], two agents from TAC/AA 2009 Competition, *AstonTAC* and *WayneAd*, were included in the experiment. The controlled experiment consisted of 40 games whose average results are shown in Table 6.

Table 6. Average results in the conducted competition

Position	Agent	Game score
1.	TacTex	57 848
2.	Mertacor	53 998
3.	Schlemazl	53 933
4.	CrocodileAgent	49 435
5.	tauagent	47 789
6.	AstonTAC	45 104
7.	EPFLAgent	44 179
8.	WayneAd	36 456

The games in the controlled experiment were configured to ensure fair capacity distribution among competing agents so each agent played ten games with high capacity, twenty games with medium capacity and ten games with low capacity. In each game, two agents had low capacity, two agents had high capacity and four of them had medium capacity. The goal of the experiment was to observe agents' behaviour in respect with the assigned capacities. The results of these observations are shown in Figure 2.

In the graph shown on Figure 2 the bars represent the ratio between the average result of a single agent in those games where the specified capacity (i.e., high, medium or low) was assigned to it and the average result of the same agent in all games. We call this ratio the intra-agent relative profitability. On the other hand, the horizontal lines represent the ratio between average results of all agents in those games where the specified capacity (i.e., high, medium and low) was assigned to them and the average results of all agents in all games. This measure represents the average intra-agent relative profitability of all agents. Finally, squares, triangles and diamonds represent the ratio between the average result of a single agent in those games where the specified capacity (i.e., high, medium or low) was assigned to it and the average result of all agents in those games where the same capacity was assigned to them. This graph enables the

[8] http://www.sics.se/tac/showagents.php

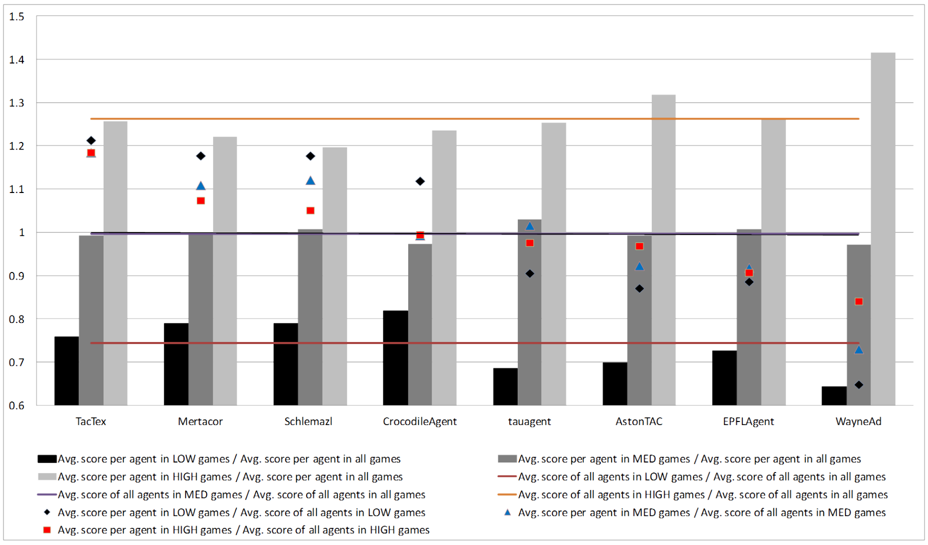

Fig. 2. Agents' relative profit with respect to assigned capacity

comparison of the single agent's average profit achieved in those games where different capacities were assigned to it with: i) its average profit in all games, and, ii) the average profit of all agents in the games where the same capacity was assigned to them. While the former measure enables us to compare the profitability of agent's strategies in those games where different capacities were assigned to it with the agent's overall profitability, the latter measure provides relative benchmarking among different agents.

If we take a look at the average values of profits achieved in the games with the assigned low capacity, we can notice that the CrocodileAgent has the highest intra-agent relative profitability among all agents in the controlled experiment. However, the CrocodileAgent has smaller intra-agent relative profitability than the average intra-agent relative profitability of all agents in the games where medium and high capacity were assigned to it. At the same time, we can also notice that, when comparing the average score of different agents, the CrocodileAgent has a 12% better score than all agents' average in the games where low capacity was assigned to them, while its performance in both medium and high capacity games is equal to all agents' average in those games.

From this analysis we can identify certain CrocodileAgent's deficiencies. The improvement of those drawbacks in future versions could significantly increase its profits. Namely, we can conclude that the CrocodileAgent should examine the possibility of using other strategies in the games when the agent is assigned with medium or high capacity in order to increase its profit in those games.

Another interesting thing we can learn from Figure 2 is that the relative intra-agent profitability of the TacTex agent, the best agent in the competition, is approximately equal to the average intra-agent profitability of all competing agents (for all three capacity allocations). Furthermore, it is attention-grabbing that the WayneAd agent, who placed last in the competition, has the highest relative intra-agent profitability in those games where the high capacity was assigned to it. However, WayneAd's weakest absolute results in medium and low capacity games are the reason for it placing last in the competition.

After analysing the impact of the assigned capacity on agents' achieved profits, we have also analysed the correlation of the achieved profit and the query category. As mentioned earlier, there are three types of queries (i.e., F0, F1 and F2) that users generate. Each advertiser selects an ad for display for each query type, choosing between a generic or targeted ad which mentions a particular product [11]. The agents' profits from all games in the controlled experiment were grouped based on the type of user query from which the transaction originated from. The mentioned distribution of agents' profits is presented in Figure 3.

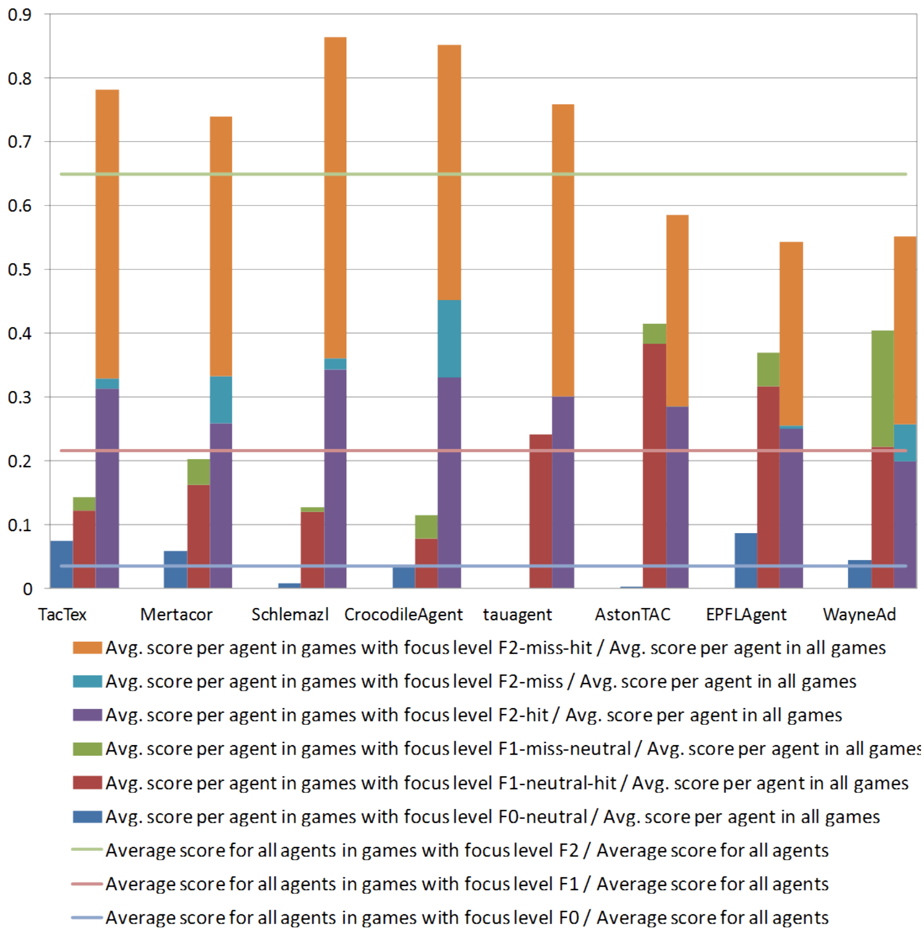

Fig. 3. Distribution of profit achieved from transactions originating from different query classes

In the graph shown on Figure 3, the horizontal lines represent the ratio between average results of all agents achieved from transactions originating from queries of the specified type (i.e., F0, F1 and F2) and the average results of all agents for all types of queries. On the other hand, the bars represent the ratio between the agent's average result achieved from transactions originating from queries of the specified type (i.e., F0, F1 and F2) and the agent's average result achieved from transactions originating from all types of queries.

If we look at the average values of profits originating from different query types, we can notice that agents, who placed higher in the competition, also achieve higher profits on targeted ads. The last three agents in the competition (i.e., AstonTAC, EPFLAgent and WayneAd) achieved the lowest relative profits from transactions originating from focus level F2 queries and the highest relative profits from transactions originating from focus level F1 queries.

Expressed in percentages and in correlation with the average results for queries of all types, the TacTex made 78% of its total profit from transactions originating from focus level F2 queries and only 22% of its total profit was made from transactions originating from focus level F0 and focus level F1 queries. Agent Schlemazl, who has the highest relative score for transactions originating from focus level F2 queries, achieved 86% of its total profit from those transactions.

Another interesting thing we can learn from Figure 3 is that overall profits are not highly correlated with the fraction of clicks received under manufacturer specialty (i.e., focus level F2-hit). Therefore, we can conclude that shifting spend towards queries focusing on manufacturer speciality is not a guarantee of greater profitability for agents, despite bonuses they get if such transactions take place.

If we analyse the CrocodileAgent's profit distribution depending on the type of user query that the transactions originated from, we can notice that the CrocodileAgent achieves approximately 85% of its total profit from transactions originating from focus level F2 queries. This percentage is higher for the CrocodileAgent than for two leading agents, TacTex and Mertacor. On the other hand, the CrocodileAgent achieved lower relative profits than the TacTex and Mertacor from transactions originating from focus level F1 queries. We can conclude that the CrocodileAgent should try to redistribute a few percent of its profit obtained from F2-queries to F1-queries, while maintaining the share of F0-queries.

6 Conclusion and future work

The Trading Agent Competition Ad Auctions (TAC/AA) game enables the academic community and advertising industry to analyse the effects of various bidding strategies by running simulations of sponsored search scenarios. Furthermore, the fact that sponsored search is the most profitable online advertising revenue format also gives great importance to ad auction research.

In this paper, we presented bidding strategies of the CrocodileAgent, the representative of University of Zagreb in the TAC/AA 2010 Tournament. Furthermore, we conducted a controlled experiment with the best-ranked agents from 2010 and 2009 TAC/AA Finals. Based on the analysis of the controlled experiment, we: i) explained some of the reasons why certain agents performed better than others, and, ii) identified CrocodileAgent's behaviours that should be improved in order to boost its performance.

For future work we plan to enhance CrocodileAgent's performance by implementing guidelines for improvements derived from the analysis of the controlled experiment. Namely, we will: i) redesign strategies for achieving profits in the games with assigned medium or high capacity to our agent, and ii) redistribute a part of the relative CrocodileAgent's profit made from transactions originating from focus level F2 queries to transactions originating from focus level F1 queries.

Acknowledgments

The authors acknowledge the support of research project "Content Delivery and Mobility of Users and Services in New Generation Networks" (036-0362027-1639), funded by the Ministry of Science, Education and Sports of the Republic of Croatia.

References

1. Feng, J., Bhargava, H.K., Pennock, D.M.: Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. INFORMS Journal on Computing **19**(1) (2007) 137–148
2. Chen, J., Feng, J., Whinston, A.B.: Keyword auctions, unit-price contracts, and the role of commitment. Production and Operations Management **19**(3) (2010) 305–321
3. Easley, D., Kleinberg, J.: Networks, Crowds, and Markets: Reasoning About a Highly Connected World. Cambridge University Press (2010)
4. Joachims, T., Granka, L., Pan, B., Hembrooke, H., Gay, G.: Accurately interpreting clickthrough data as implicit feedback. In: Proceedings of the 28th ACM SIGIR conference on Research and development in information retrieval. SIGIR '05, New York, NY, USA, ACM (2005) 154–161
5. Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: Proceedings of the international conference on Web search and web data mining. WSDM '08, New York, NY, USA, ACM (2008) 87–94
6. Jeziorski, P., Segal, I.: What makes them click: Empirical analysis of consumer demand for search advertising. Economics Working Paper Archive 569, The Johns Hopkins University, Department of Economics (2010)
7. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Proceedings of the 7th ACM conference on Electronic commerce. EC '06, New York, NY, USA, ACM (2006) 218–227
8. Lahaie, S., Parkes, D.C., Pennock, D.M.: An expressive auction design for online display advertising. In: Proceedings of the 23rd national conference on Artificial intelligence - Volume 1, AAAI Press (2008) 108–113
9. Varian, H.R.: Online ad auctions. American Economic Review **99**(2) (2009) 430–34
10. Acharya, S., Krishnamurthy, P., Deshpande, K., Yan, T., Chang, C.C.: A Simulation Framework for Evaluating Designs for Sponsored Search Markets. In: 16th International World Wide Web Conference. (2007)
11. Jordan, P.R., Cassell, B., Callender, L.F., Wellman, M.P.: The ad auctions game for the 2009 trading agent competition. Technical report (2009)

12. Jordan, P.R., Wellman, M.P.: Designing an ad auctions game for the trading agent competition. In: Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets. Volume 59 of Lecture Notes in Business Information Processing. Springer Berlin Heidelberg (2010) 147–162
13. Arunachalam, R., Sadeh, N.M.: The supply chain trading agent competition. Electronic Commerce Research and Applications **4**(1) (2005) 66–84
14. Podobnik, V., Petric, A., Jezic, G.: An Agent-Based Solution for Dynamic Supply Chain Management. Journal of Universal Computer Science **14**(7) (2008) 1080–1104
15. Sardinha, A., Benisch, M., Sadeh, N., Ravichandran, R., Podobnik, V., Stan, M.: The 2007 procurement challenge: A competition to evaluate mixed procurement strategies. Electronic Commerce Research and Applications **8**(2) (2009) 106–114
16. Niu, J., Cai, K., Parsons, S., Gerding, E., McBurney, P.: Characterizing effective auction mechanisms: insights from the 2007 TAC market design competition. In: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems, AAMAS'08, International Foundation for Autonomous Agents and Multiagent Systems (2008) 1079–1086
17. Petric, A., Podobnik, V., Grguric, A., Zemljic, M.: Designing an Effective E-Market: An Overview of the CAT Agent. In: Proceedings of AAAI-08 Workshop on Trading Agent Design and Analysis, TADA-08, AAAI Press (2008) 62–65
18. Block, C., Collins, J., Ketter, W.: Agent-based competitive simulation: Exploring future retail energy markets. In: Twelfth International Conference on Electronic Commerce (ICEC 2010), ACM (2010) 67–76
19. Lahaie, S., Pennock, D.M.: Revenue analysis of a family of ranking rules for keyword auctions. In: Proceedings of the 8th ACM conference on Electronic commerce. EC '07, New York, NY, USA, ACM (2007) 50–56
20. Berg, J., Greenwald, A., Naroditskiy, V., Sodomka, E.: A first approach to autonomous bidding in ad auctions. In: EC 2010 Workshop on Trading Agent Design and Analysis, TADA 2010, New York, NY, USA, ACM (2010)
21. Pardoe, D., Stone, P.: A particle filter for bid estimation in ad auctions with periodic ranking observations. In: EC 2010 Workshop on Trading Agent Design and Analysis, TADA 2010, Cambridge, Massachusetts, ACM (2010)
22. Cigler, L.: Semester project: Bidding agent for advertisement auctions. Technical report, Ecole Polytechnique Federale de Lausanne, Ecole Polytechnique Federale de Lausanne (2009)
23. Witten, I.H., Frank, E.: Data Mining: Practical Machine Learning Tools and Techniques. Morgan Kaufmann Publishers, San Francisco, CA, USA (2005)