



TECHNICAL UNIVERSITY OF CRETE

MULTIAGENT SYSTEMS
FINAL REPORT

TAC AA Agent Description

Nikolaos Pavlakis, Ioakim Perros, Eleftherios Soulas

November 1, 2012

Contents

1	Introduction	1
2	Tac Ad Auctions Game	2
2.1	Advertisers	2
2.2	Publishers	3
2.3	Users	3
3	Background knowledge	3
4	Our approach	4
4.1	Bidding strategy	4
4.2	Ad selection	8
4.3	Spending limits	8
5	Results	8
5.1	Agent’s performance	8
5.2	Useful conclusions	8
5.3	Future work	9
A	Appendix	11

1 Introduction

Sponsored search auctions constitutes the most important source of revenue for search engine companies, offering new opportunities for advertisers. The Trading Agent Competition (TAC) Ad Auctions tournament is one of the first attempts to study the competition among advertisers for their placement in sponsored positions along with organic search engine results. In this report, we describe our agent "SparTUCore" (a.k.a. "TUCTAC"), a simulation-based game theoretic agent that successfully competed in the 2012 Tac Ad Auctions tournament ,which took place in Valencia (4-5 June), in conjunction with AAMAS-12 and the workshop on Trading Agent Design and Analysis.

Section 2 provides a brief description of the game. A compact review of the papers we read is presented at section 3. Section 4 builds the background upon which our agent was based and gives an analysis of our approach. Finally, a discussion of our agent’s performance at both the official tournament and the internal course competition, is given at section 5, coupled with conclusions and thoughts regarding future work.

2 Tac Ad Auctions Game

Sponsored search auctions are open, highly complex mechanisms, that are non-dominant-strategy-solvable, hence bidding strategies are a topic of active research. To investigate their behavior, a realistic agent-based simulator seems essential. The Ad Auctions (AA) platform in the international Trading Agent Competition (TAC) is such a system. The TAC AA game specifications are defined in detail in [1]. To familiarize the reader with the game, we will provide some basic information about the entities involved and the interactions between them.

In TAC AA tournament, there are three main types of entities, the publisher, a population of 90000 users, and eight advertiser entrants represented by autonomous software agents. The advertisers compete against each other for advertisement placement, across search pages. Each one of the search pages contains search engine results for one of the queries of 16 different keyword sets. In order to promote their products, the agents participate in ad auctions by submitting a bid and an ad to the publisher for the query (set of keywords) they are interested in. Ads are ranked on each search page, based on a generalized method that interpolates between rank-by-bid and rank-by-revenue schemes. Each day, users, according to their preferences and state, remain idle, search, click on ads and make purchases (conversions) from the advertisers' websites. The products being traded are combinations of three brands and three types of components from the domain of home entertainment. The small number of products enables competing teams to focus only on a small set of predefined keywords, abstracting away from the problems of keyword selection. The three manufacturers (namely, Lioneer, PG and Flat) and the three types of devices (TV, Audio and DVD) constitute a total of nine products. The simulation runs over 60 virtual days, with each day lasting 10 seconds [1].

2.1 Advertisers

Each advertiser is a retailer of home entertainment products and can supply the user with any of the nine products available. Upon initialization of each simulation, advertisers are given a component and a manufacturer specialty, yielding an increase in conversion rates for the former and an increase in profit per unit sold for the later - which will prove to be crucial amounts for our agent's performance. Additionally, entrants are assigned a weekly maximum stock capacity $C^{Cap} \in \{C^{LOW}, C^{MED}, C^{HIGH}\}$, so conversions above this threshold are less likely to happen during this week (5 working days).

2.2 Publishers

The publisher runs a Generalized Second Price auction (its most important aspect is that the highest bidder wins, but the price paid is the second-highest bid) to determine the rank of bids and determine the payment per click. The ad placement algorithm takes into account predefined reserve scores. There is a reserve score below which an ad will not be posted, and one above which, an ad will be promoted. If the spending limit set by an agent is passed, the rankings are recalculated. The auction implemented is a GSP, where the ranking takes into account the quality of the advertisements, weighted by a squashing parameter that is disclosed to the entrants at the beginning of the game. [1]

2.3 Users

Each user has a unique product preference and can be in different states representing his or her searching and buying behavior (i.e. non-searching, informational searching, shopping, with distinct levels of shopping focus, and transacted). The product preference distribution is even for all products. Users submit three kinds of queries, defined by their focus level for a total of 16 queries. There is one (1) F0 query, where no manufacturer or component preference is revealed, six (6) F1 queries, where only the manufacturer or the product type is included in the query and nine (9) F2 queries, where the full product definition (manufacturer and type) is exposed. Users' daily state transition is modeled as a Markov chain. Non-searching and transacted agents do not submit queries. Informational agents submit one of the three queries by selecting any one of the them uniformly and focused users submit a query depending on their focus level. While both information seeking and focused users could click on an ad, only focused users make purchases and go to the transacted state. After clicking on an ad, whether a conversion will be made or not depends on users state, advertisers' specialty and remaining distribution capacity. If the reader seeks for a detailed analysis of the game specifications, useful citations are [1], [2].

3 Background knowledge

In this section, we refer to related work, which includes agent strategies from pre-existing competitive agents. The majority of strategies is focused on two target metrics, namely the *Return on Investment*, i.e., the ratio of profit to cost and the *Value per Click*, i.e. the expected profit from a conversion given

a click, and combined with multiple choice knapsack (MCKP) models to deal with the distribution constraint effect.

The consecutive winner of past competitions, TacTex [3], implements a strategy in which the focus lies in estimating unknown values and distributions such as value per click estimation, user estimation, using a particle filtering technique. Another competitive participant, EPFLAgent [5] exploits the observation that focusing on queries matching the agent’s manufacturer specialty, results in greater profits. This approach is also adopted by our agent, as a component of our bidding strategy.

Vorobeychik, at his work [9], provides a general bidding strategy, estimating (for each query per day) each bid with the formula: $bid = \alpha \cdot VPC$, which constituted our main approach concerning the bidding estimation. Mertacor [7], winner of Tac Ad Auctions 2012, besides implementing the same bidding estimation strategy with the aforementioned work of Vorobeychik [9], estimates the distribution constraint effect with an interesting approach, which we embedded in our strategy. AstonTAC [8] estimates the market VPC, which is the VPC minus relevant cost, and then bids a proportion of this value based on the critical capacity (i.e. capacity beyond which the expected cost is higher than the corresponding revenue) and the estimated quality factor for the advertiser. Priority is also given to the most profitable queries.

4 Our approach

In order to compete in a Tac Ad Auctions competition, as mentioned earlier, an agent has to determine three aspects of the game. Our approach ¹ considering each one of them, is described below:

4.1 Bidding strategy

A two-stage bidding process is adopted due to the fact that for the first two days, we have no reports available concerning statistics about the market. For this reason, we would prefer a static/standard way of bidding for the first phase, and a dynamic one, based upon market reports, for the second one.

1. First two days

At this game’s phase, the formula used (based on [8]) is as follows:

$$b_0 = V_{\text{static}} H_q H_c$$

¹It is worth mentioning that our code’s skeleton is based on Spartac’s agent layout (which was provided by the instructor).

The target of the above expression is to increase/decrease a static amount of bid, by taking into consideration two main thoughts: first, that we want to favour the products of our manufacturer's specialty (in order to gain the respective bonus) and also that we want to be more aggressive considering our bids if our products' capacity is high and more passive if this capacity is low. As a result, H_q is introduced in order to counterbalance the possible increase of conversion from high-value queries under restricted total conversion, because relatively low bid prices should be set for relatively low-value queries. Likewise, H_c is used so that the agent is more conservative when bidding under low capacity and more offensive when bidding under high capacity. More specifically our choices for these two parameters follow:

$$H_q = \begin{cases} 1.2, & \text{if } q.getManufacturer() = \text{manufacturerSpecialty} \\ 0.8, & \text{if } q.getComponent = \text{null} \ \&\& \ q.getManufacturer() = \text{null} \\ 1, & \text{otherwise} \end{cases}$$

$$H_c = \begin{cases} 0.85, & \text{if } \text{distributionCapacity} = \text{CLOW} \\ 1, & \text{if } \text{distributionCapacity} = \text{CMED} \\ 1.15, & \text{if } \text{distributionCapacity} = \text{CHIGH} \end{cases}$$

V_{static} is based upon this formula:

$$V_{\text{static}} = \overbrace{\underbrace{P_{\text{conversion}}}_{\text{Conv. rate}} \cdot \underbrace{V_{\text{con}}}_{\text{Conv. value}}}_{\text{expected revenue-per-click}} \cdot r_{\text{discount}}$$

The intuition behind the above formula lies into calculating the expected revenue-per-click through multiplying conversion rate ($P_{\text{conversion}}$) by conversion value (V_{con}), and also by r_{discount} in order to obtain the Market-based Value Per Click (MVPC) of each query.

Furthermore, $P_{\text{conversion}}$ and V_{con} are described through:

$$P_{\text{conversion}} = \begin{cases} \frac{(1+CSB)\pi_l}{1+CSB\pi_l}, & \text{if } q_c = \text{CS} \\ \frac{2}{3}\pi_l + \frac{1}{3}\frac{(1+CSB)\pi_l}{1+CSB\pi_l}, & \text{if } q_c = \text{null} \\ \pi_l, & \text{otherwise} \end{cases}$$

$$V_{\text{con}} = \begin{cases} USP(1 + MSB), & \text{if } q_m = \text{MS} \\ \frac{2}{3}USP + \frac{1}{3}USP(1 + MSB), & \text{if } q_m = \text{null} \\ USP, & \text{otherwise} \end{cases}$$

It is worth mentioning that the cases where $q_c = \text{null}$ and $q_m = \text{null}$ refer to queries where the chance of a matching component/manufacturer specialty is $\frac{1}{3}$, so the respective formulas take that amount into consideration for the $\frac{1}{3}$ of the cases, and the π_l/USP amounts for the remaining $\frac{2}{3}$.

2. Third day & beyond

The second part of our strategy is based upon Vorobeychik's [9], bidding tactic, where each query's bid is calculated by the following equation:

$$\boxed{bid = \alpha \cdot VpC}$$

so the whole bidding tactic lies into estimating α and Value-per-Click.

In [7], they employ a sophisticated approach in order to determine the optimal α , by turning the problem into an associative n-armed bandit problem (as in [6]). Despite this, the authors in [7], based on experiments, concluded that setting α to a value of 0.3 is best response for more reasonable strategies and further accuracy in estimating α , does not significantly boost agent's performance.

By diving into the process of estimating Value-per-Click, we tried to simplify the approach followed by [7] and expressed this amount by the following formula:

$$VpC = Pr\{\text{conversion}\} \cdot E\{\text{revenue}\}$$

More specifically, each one of the above amounts, is given by estimating next day's:

$$Pr\{conversion\} = \begin{cases} \eta(\pi_l I_d, 1 + CSB), & \text{if user matches CS (Component Specialty)} \\ \frac{2\pi_l I_d}{3} + \frac{\eta(\pi_l I_d, 1 + CSB)}{3}, & \text{if user doesn't specify component} \\ \pi_l I_d, & \text{otherwise} \end{cases}$$

$$E\{revenue\} = \begin{cases} \frac{(USP(3+MSB))}{3}, & \text{if MS not defined in q} \\ USP \cdot (1 + MSB), & \text{if MS matched in q} \\ USP, & \text{if MS not matched in q} \end{cases}$$

Both of the above formulas were mainly extracted from the specifications' manual [1]. A similar approach to the formulas calculating $P_{conversion}$ and V_{con} , is followed for the case when query' s component/manufacturer is omitted.

The crucial factor of estimating the amount of $P_{conversion}$ is the distribution constraint effect (I_d), which is again defined by the specification manual [1], in the following way:

$$I_d = \lambda^{((\sum_{i=d}^{d-(W-1)} c_i) - C^{Cap})^+}$$

As noticed in [10], the distribution constraint effect, I_d , is the second most influential factor in an advertiser's performance after the manufacturer specialty, affecting radically the probability of purchase.

Based on [7], we decided to calculate and incorporate I_{d+1} to the above formulas, instead of I_d , in order to make a sort of prediction for the distribution constraint effect for day d+1.

The difficulty in calculating this amount, lies into the fact of estimating the values of c_i and c_{i+1} , which indicate our future conversions for the next two days. As for the value of c_i , it is the mean of the former three days' conversions, when our capacity is limited to 300 and to one quarter of their sum, otherwise. On the other hand, to be more conservative, we further decrease (e.g. by half of c_i) the conversions predicted for the last day in our window. In section 5, we will further elaborate on the success of the above choices.

4.2 Ad selection

The way we choose the kind of ads to show to users is determined by:

$$Ad = \begin{cases} generic, & \text{if } q.getManufacturer()=null \ \& \\ & q.getComponent()=null \\ Ad(MS, q.getComponent()), & \text{if } q.getManufacturer()=null \\ Ad(q.getManufacturer(), CS), & \text{if } q.getComponent()=null \\ Ad(MS, CS), & \text{otherwise} \end{cases}$$

The intuition behind this, is the forwarding of targeted ads of our component/manufacturer specialty, in cases when component/manufacturer slots are not defined by the query. As a result, with probability of $\frac{1}{3}$, we gain improved Click-Through-Rate/Revenue. In section 5, we will discuss on the success of this "gambling" technique.

4.3 Spending limits

We did not make use of daily spending limits, because we would not prefer to limit the dynamics of our main strategy. More details will be provided in the next section. In addition, we experimentally imposed campaign spending limits to 3500 for the lowest possible capacity and to 4500 for the case when C_{cap} equals 450.

5 Results

5.1 Agent's performance

If we would like to discuss about the results of our agent's games, it is worth mentioning that our agent won the internal course (CSC-517) competition and also qualified for the Tac Ad Auctions 2012 finals, where we were placed 6th out of 8 participants.

The qualified agent for the tournament semifinals was created by our colleagues Aggelos Aggelidakis and Nikolaos Sapountzis - which was beaten by our agent at an internal competition we organized, in order to determine the agent that would represent our institution at competition finals.

5.2 Useful conclusions

As conclusions of our remarks from the games we played during both the debugging and the tournament phase, we would like to refer to the following:

1. Our strategy relies on spending large amounts in a narrow period of time (1-2 days), and expecting large profits from these sales, while remaining idle until there is a capacity renewal. We believe this experimental remark appears to be useful for the understanding of the following results.
2. The influence of possessing the lowest possible capacity is crucial to our agent's performance. In this case, we noticed that it was essential to set a campaign spending limit (as well as in the case of having $C_{\text{cap}} = 450$), in order to avoid spending our profits. In the case where we would not set a spending limit, we would end up getting a lot of clicks, but not as many conversions, due to the fact that the capacity would be limited and the probability of conversion would be low.
3. A very important aspect of successful performance is the maintenance of capacity used to a percentage of 125%. This provides us with a satisfying conversion probability (which is proportional to capacity available), leading to high profits.
4. We noticed that it is preferable for an agent to bid on queries complying with their manufacturer specialty, as opposed to bidding high on their component specialty queries, as in the former case the agent gains 40% increased profit, while the latter results in an augmented (60%) conversion probability. Thus, if we are to sell a given amount of products, we prefer that these products match our manufacturer specialty - so that we exploit the 40% aforementioned bonus.
5. We experienced improved performance if "dummy" agents were absent from a game. This type of agent reacts at times as a "saboteur" for our agent. This happens because they make, once in a while, higher bids to many auctions, even if this strategy is not rational for them. Thus, the rest of a game's agents, tend to lose profit for no real reason.
6. With respect to other competitors at TAC Ad Auction Finals, our agent was lacking a more sophisticated and algorithmically robust procedure, which would exploit all report statistics to further estimate and better predict the state of the game.

5.3 Future work

This was the first time our university participated in this worldwide competition, but we would like to continue competing in the years to come. Our first thoughts considering possible improvements of our agent consist of:

- We have to employ a dynamic spending limit calculation strategy, in order to avoid the current static one, which is inflexible.
- We aim to create a user population estimation at each state, as in [3], which could be useful in various aspects of our agent’s strategy.

References

- [1] The Ad Auctions Game **Patrick R. Jordan, Akshat Kaul, and Michael P. Wellman**
- [2] Designing an Ad Auctions Game for the Trading Agent Competition **Patrick R. Jordan and Michael P. Wellman**
- [3] TacTex09: Champion of the First Trading Agent Competition on Ad Auctions **David Pardoe, Doran Chakraborty, and Peter Stone**
- [4] A Particle Filter for Bid Estimation in Ad Auctions with Periodic Ranking Observations **David Pardoe and Peter Stone**
- [5] EPFLAgent: A Bidding Strategy for TAC/AA. **Ludek Cigler and Elias Mirza** Presentation in Workshop on Trading Agent Design and Analysis (TADA), ACM EC 2010, July 2010
- [6] Reinforcement Learning: An Introduction. **Richard S. Sutton and Andrew G. Barto**. MIT Press, Cambridge, MA, 1998.
- [7] An Adaptive Proportional Value-per-Click Agent for Bidding in Ad Auctions **Kyriakos C. Chatzidimitriou, Lampros C. Stavrogiannis, Andreas L. Symeonidis and Pericles A. Mitkas** The 2011 Workshop on Trading Agents and Design Analysis
- [8] Designing a Successful Adaptive Agent for TAC Ad Auction **Meng Chang and Minghua He and Xudong Luo**
- [9] Simulation-Based Game Theoretic Analysis of Keyword Auctions with Low-Dimensional Bidding Strategies **Yevgeniy Vorobeychik**
- [10] Practical Strategic Reasoning with Applications in Market Games. Patrick R. Jordan. PhD thesis, University of Michigan, Ann Arbor, MI, 2010.

A Appendix

Position	Agent	Average Score	Average Score - Zero	Games Played	Zero Games
1	tau (tacaal , taca2)	54 673	54 673	64	0
2	Schlemazl (tacaal , taca2)	54 182	54 182	64	0
3	Mertacor (tacaal , taca2)	52 099	52 099	64	0
4	epflagent (tacaal , taca2)	48 601	48 601	64	0
5	Time (tacaal , taca2)	47 397	47 397	64	0
6	TUCTAC (tacaal , taca2)	43 719	43 719	64	0
7	fsr (tacaal , taca2)	43 206	43 206	64	0
8	CrocodileAgent (tacaal , taca2)	41 015	41 015	64	0
9	THA (tacaal , taca2)	33 967	34 506	64	1
10	WayneAd (tacaal , taca2)	30 646	30 646	64	0

Figure 1: Agent’s performance at Tac Ad Auctions 2012 semifinals

Position	Agent	Average Score	Average Score - Zero	Games Played	Zero Games
1	Mertacor (tacaal , taca2)	52 960	52 960	64	0
2	Schlemazl (tacaal , taca2)	52 701	52 701	64	0
3	tau (tacaal , taca2)	51 913	51 913	64	0
4	epflagent (tacaal , taca2)	47 240	47 240	64	0
5	Time (tacaal , taca2)	46 996	46 996	64	0
6	TUCTAC (tacaal , taca2)	45 594	45 594	64	0
7	CrocodileAgent (tacaal , taca2)	41 990	41 990	64	0
8	fsr (tacaal , taca2)	40 916	40 916	64	0

Figure 2: Agent’s performance at Tac Ad Auctions 2012 finals

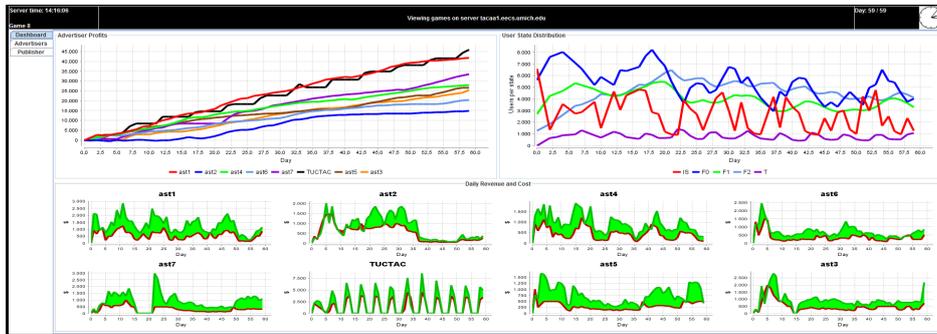


Figure 3: Example of simulation results of our agent's performance (TUC-TAC)



Figure 4: Example of simulation results of our agent's performance (TUC-TAC). It is depicted that in a game when the total profit reached almost 60k, the capacity used indication lies between 125 and 175 %.



Figure 5: Example of simulation results of our agent’s performance (TUC-TAC). It is depicted that in a game when the total profit reached almost 35k (10k lower than our average), the capacity used indication is totally unstable and even reaches values of 250 %. The difference between the previous and the current figure at this indication, seems to justify the dissimilarity at the total profit, caused by the lack of capacity available in the current case. This is why we aim to employ a more sophisticated strategy in the case of possessing the lowest available capacity.