# Truthful polynomial time optimal welfare keyword auctions with budget constraints

Elan Pavlov*

## ABSTRACT

In keyword auctions, advertisers bid for keywords. Currently, bids and allocations are made on a per keyword basis. We show that this mechanism is neither efficient nor truthful under demand or budget constraints and provide a truthful efficient (quasi) polynomial mechanism for allocating keywords under budget constraints. As a side effect our proposed mechanism has higher revenue than the current mechanism in use by search engines.

## 1. INTRODUCTION AND RELATED WORK

There has been much work (e.g., [17, 8]) done on keyword auctions. In a general keyword auction advertisers bid on a set of keywords (sometimes with a budget constraint as in [17, 19, 15]) and a search engine allocates each keyword to the highest bidder and charges the winner a "generalized second price" for the auction consisting only of that keyword. There are also papers which charge a more sophisticated price which is truthful [1] but again only for the auction consisting of a single keyword.

Multiple keywords are treated independently and allocations are made independently for each keyword. Although some authors (e.g., [17, 19, 15]) utilize budget constraints which can be a result of previous keywords and/or limit the amount of budget allocated for each keyword in general the *allocation* is not dependent on multiple keywords.

This is somewhat puzzling. Keyword auctions can be seen as a special case of a combinatorial auction and since keywords are (partial) substitutes the optimal allocation can not be made locally but must be chosen across the entire possibility of allocation with complex interdependencies between keywords.

It is true that when there are no budget constraints the optimal allocation can be chosen by allocating the advertiser with the highest *payoff* (value * click through rate) for each impression but this is no longer the case when budget constraints are allowed (as is the case in theory and even in practice).

In fact, the keyword auction with multiple keywords can be considered a special case of a *combinatorial auction (CA)*. In this case a budget constraint can be understood as modeling a risk neutral[1] advertiser who enters a multi-unit auction for click-throughs and is interested in a finite number of click-throughs. Of course, it would be easier to explicitly state the number of click-throughs desired but a budget is a good proxy for this when lacking complete information on the costs. In this paper we will assume budget constraints are given in terms of number of desired clickthroughs. By interpreting budget constraints as a limit on the desired number of clickthrough we circumvent [5] and hence it is possible to maximize welfare in this model. In section 7 we sketch a way of translating monetary bounds on the budgets into desired number of click-throughs in quasi-polynomial time.

Combinatorial auctions are a widely studied mechanism (see e.g., [6]). The main problem with looking at the keyword auction as a special case of a CA is that it is impossible to approximate the welfare in a CA beyond a square-root factor in polynomial time [16]. However, there are known special cases in which the welfare can be approximated (e.g., [7, 18, 16]) or even achieved optimally for special domains (e.g., [11, 10]) possibly using alternative solution concepts [3, 2].

Our goal is to look at keyword auctions as a special case of CA with multiple keywords acting as substitutes. In order to circumvent the known hardness results we show that keywords auctions are a special case of CA which can be solved in (quasi) polynomial time.

The main technical tool which we use to avoid the hardness results on CAs is a *generalized min cost flow algorithm* [20]. By phrasing our problem as a flow problem we can maximize the welfare. Our setting in which we can compare possible allocations is a natural one for employing flow algorithms.

Our proposed mechanism allows users to declare a desired number of click-throughs (which in most of this paper except section 7 stands in lieu of the budget) and the mechanism will allocate impressions so as to maximize the total social welfare. Our mechanism assumes that the probabilities of click-throughs for each advertiser is known to the mechanism but this is the generally made assumption. If the probabilities are unknown the mechanism can learn them via such results as [12, 4].

*E15-313 Media Lab, MIT, Cambridge MA, 02149. Email: elan@mit.edu

[1] Obviously for risk-negative advertisers the budget constraint is easily understood.

## 2. MODEL

We assume that there are $n$ advertisers and $m$ possible keywords. For each advertiser $i$ and keyword $j$ there is a probability $p_{ij}$ that a user searching for keyword $j$ will click on the ad of advertiser $i$[2]. Note that if $p_{ij}$ was independent of $j$ then all of the keywords would be the same and hence the problem would be reduced to a multi-unit auction. We assume that the probabilities of the CTRs $p_{ij} \in \mathcal{Q}$[3]. Our algorithms will be quasi-polynomial and we denote by $B$ the largest integer that appears in either the numerator or denominator of $p_{ij}$.

We assume that advertisers are risk neutral, with quasi-linear utility and maximize their own utility.

We assume that the mechanism knows all of the $p_{ij}$. In practice, this commonly made assumption is problematic. However, by using a a learning mechanism or heuristics this assumption is not completely groundless in practice. In most work in the field (except for [12]) this is a commonly made assumption.

Each advertiser $i$ has two parameters which are the value per click-through $v_i$ and the total number of click-throughs desired $n_i$. We assume that $v_i, n_i$ are *private information* and known only to advertiser $i$.

For each possible keyword $j$ the number of appearances $a_j$ of keyword $j$ is known to the mechanism. This can actually be justified based on the law of large numbers and the vast amount of data collected by the search engines. In the conclusions we deal with the case where this is unknown.

It can be seen that this is a special case of a multi-unit combinatorial auction and hence merely using the standard VCG approach is not guaranteed to terminate in polynomial time (nor to approximate it better than a square-root factor in polynomial time).

The goal of our mechanism is to allocate slots to advertisers so as to maximize the social welfare. Furthermore, this should be *individually rational* as well as *incentive compatible.*

Since we look at a probabilistic setting (where click-throughs are a random event) we choose to maximize the expected welfare. This means that we allocate all of the possible impressions and the start of the protocol and do not take into account actual clicks throughs. It is possible although difficult to deal with actual events assuming that the advertisers have no private information about any a-priori distribution.

Since we will compare the revenue and welfare of our proposed mechanism to the mechanism which is currently used for allocating keywords by the search engines[4] we denote our proposed mechanism by $\mathcal{A}$ and the current mechanism used by the search engines by $\mathcal{A}'$.

It is interesting to compare our result with [9]. In [9] it is shown that a linear fraction of the optimal value can be achieved by advertisers bidding a fixed value. This however assumes that the current mechanism used by search engines[5] (which loses a lot of welfare) is in effect. [9] does not compare the loss of welfare incurred by the choice of mechanism.

Another interesting comparison is with [17]. [17] discover an approximation of the welfare by allocating at each time period based on the remaining budget. Unfortunately, this is not truthful and in fact advertisers can arbitrarily increase their utility by lying. However, our result makes the strong assumption that there is a known distribution on the appearance of keywords and our result only maximizes the expected welfare. Furthermore, [17] looks at the case when all probabilities are the same and hence the problem is one of determining a matching between bidders and keywords and does not take into account possibly different probabilities.

[2]It is also possible to generalize our mechanism to the case where players have different values for different keywords.

[3]We assume that this is the smallest representation. This assumption is necessary for our solution to complete in polynomial time.

[4]Ranking by payoffs and charging second price

[5]The current mechanism ranks bidders by payoff and charges critical values. We assume for ease of presentation that there is a single slot associated with each keyword. The generalization to multiple slots with different quality is trivial.

## 3. THE PROBLEM

The problem with the current mechanism for allocating keywords arises is when a closely competed (i.e. is desired by many advertisers) slot "runs out" and the difference in CTRs (and hence payoffs) for the other slots are big. The following example shows that the current mechanism is neither efficient nor truthful.

1. *There are three advertisers* $1, 2, 3$ *and two keywords* $1, 2$. *The probabilities are as follows:* $p_{11} = 1, p_{12} = 1 - \epsilon, p_{21} = 1 - \epsilon, p_{22} = 0, p_{31} = p_{32} = \epsilon' > \epsilon$ *and the valuations are* $v_1 = v_2 = v_3 = 1$ *budgets are* $b_1 = b_2 = b_3 = 1$. *There is also a single time slot for each keyword.*
*If advertisers compete independently for each keyword that advertiser* $1$ *will win keyword* $1$ *and advertiser* $3$ *will win keyword* $2$ *(since advertiser* $1$ *doesn't have sufficient budget for both keywords). The welfare of this allocation is* $v_1 * p_{11} + v_3 * p_{32} = 1 - \epsilon + \epsilon' = 1$
*However, the optimal allocation is to allocate player* $2$ *keyword* $1$ *and allocated player* $1$ *keyword* $2$. *The welfare in this case is* $v_2 * p_{21} + v_1 * p_{12} = 2 - 2\epsilon = 2$

Note that although the example has a single impression for each keyword, it is possible to build a similar example with multiple impressions. It can be shown via a charging argument that the loss in welfare can be bounded by a factor 2 and the above example shows that this bound is tight. However, the example must assume budget constraints:

REMARK 3.1. *If there is no budget constraint the efficient allocation is to allocate the keyword to the highest bidder for that keyword.*

This is due to the fact that if there is a positive marginal utility for each impression than an advertiser will want to win as many slots as possible. Maximizing welfare then demands that we allow an advertiser to win as many keywords as possible which in turn can be achieved by running an auction for each keyword separately.

REMARK 3.2. *This example also shows that the current mechanism is not truthful inasmuch as player 1 has an incentive to lie about his value for keyword 1 in order to reduce payments. Even if players are restricted to a single value which is independent of the keyword (which is not the case today) the incentive for player 1 to lie remains.*

To deal with the problem of having competitive keywords sell out while we still have excess "capability" in differing slots we add the ability to "move" advertisers between keywords. This will be done when such movement will improve the total welfare even if this results in some advertiser not receiving the keyword with the highest probability. However, we will ensure that the advertiser is not harmed by being moved to another keyword (by allocating more impressions).

We cast our problem as a generalized minimum cost flow problem which can be solved in quasi-polynomial time. In the next section we review the results from generalized minimum cost flow problem which we use.

# 4. GENERALIZED MINIMUM COST FLOW PROTOCOLS

In this section we briefly recap the definitions and results for generalized min-cost flow (GMCF). We utilize the notations and results of [20].

Given a graph $G = (V, E, u, c, \gamma)$. $V$ is the vertex set, $E$ is the directed edge set. $u : E \to \mathcal{R}_+$ is a capacity function, $c : E \to \mathcal{R}$ is a cost function and each $e \in E$ has a positive multiplier $\gamma(e)$ called a *gain factor* associated with it. For each unit of flow entering the edge $e$ there are $\gamma(e)$ units of flow that exit the edge.

A feasible generalized circulation is a nonnegative function $g : E \to \mathcal{R}_+$ that satisfies the *flow conservation* constraints:

$$\forall v \in V \sum_{(v,w)\in E} g(v,w) = \sum_{(w,v)\in E} g(w,v)\gamma(w,v)$$

as well as the capacity constraint:

$$\forall (v,w) \in E : g(v,w) \leq u(v,w).$$

The generalized minimum cost flow problem is to find a feasible generalized flow of minimum cost. The following result is taken from [20]:

1. **(Theorem 8 in [20])** *The scaling algorithm computes an $\epsilon$-optimal generalized minimum cost circulation in* $\tilde{O}(|E|^2|V|^2 \log \frac{1}{\epsilon})$. *In*
$\tilde{O}(|E|^3|V|^2 \log B)$ *time it computes an optimal flow.*

Now that we have reviewed the result we need, we proceed to cast the question of an optimal welfare keyword auction as a generalized minimum cost flow problem.

# 5. CONSTRUCTING THE FLOW GRAPH

Assuming that all of the probabilities $p_{i,j}$ are known we build the following graph (See figure 5).

1. We construct a single source node $s$
2. We construct $n$ nodes corresponding to advertisers. Each advertiser node $i$[6] has an edge from the source $s$.

   This edge has a capacity of $n_i$ (the number of desired click throughs), a unit-cost of 0 and a scaling factor of 1.
3. We construct $m$ nodes which correspond to keywords. Between each advertiser $i$ and keyword $j$ we construct an edge with capacity $\infty$ and unit-cost of $-v_i$ the scaling factor of an edge between advertiser $i$ and keyword $j$ is $\frac{1}{p_{ij}}$.
4. Finally we construct a target node $t$ and an edge between each keyword and the target with capacity $a_j$ (the number of appearances of $j$), cost of 0 and scaling factor of 1.

[6]by abuse of notation we identify nodes with advertisers and keywords

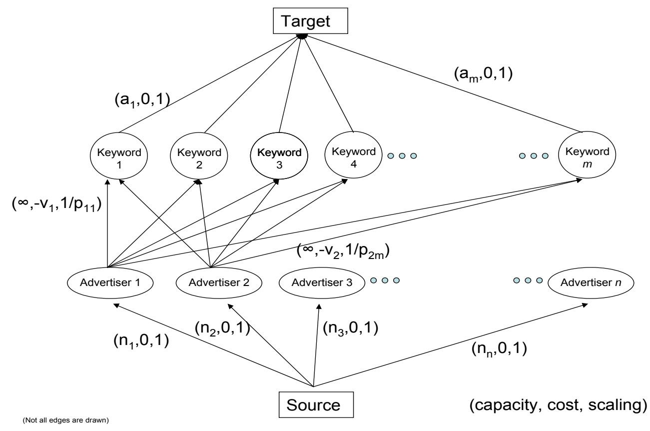


**Figure 1: A schematic view of the flow graph**

This graph is obviously a generalized flow graph with costs and as such can be solved in quasi polynomial time via [20] to find the minimum cost flow (MCF) $g : E \to \mathcal{R}_+$. We must show that $g$ induces an optimal expected welfare for the keywords problem. We also show that we can induce prices (in polynomial time) that support this flow such that advertiser $i$ will report truthfully the value of $n_i, v_i$.

We first show how to translate the MCF into an allocation. For each keyword we must decide how to allocate the keyword whenever it appears. From keyword $j$ there is an edge to the target node $t$. For a MCF $g$ if the edge $(i, j)$ from advertiser $i$ to keyword $j$ has flow $g(i, j)$ we say that $i$ contributes $\frac{g(i,j)}{p_{ij}}$ to the flow along edge $(j, t)$.

Given an MCF $g$ for any appearance of a keyword $j$ we allocate the impression of that keyword to advertiser $i$ with probability $\frac{g(i,j)}{g(j,t)p_{ij}}$. We first show that this is a probability distribution:

5.1. *The assignment which allocates to advertiser i the keyword j with probability* $\frac{g(i,j)}{p_{ij}g(j,t)}$ *is a probability distribution.*

PROOF. The proof is algebraic and is omitted. □

We call this allocation the *allocation induced by the MCF.*

We now proceed to prove that this flow has an optimal expected welfare:

2. *The allocation A induced by the MCF has an optimal expected welfare if the advertisers report their values correctly.*

PROOF. In order to show that the induced allocation has an optimal expected welfare we must show that it is indeed an allocation.

5.2. *The induced allocation maintains the constraints of the problem.*

Proof. By the definition of the graph and the fact that we have a feasible flow, no advertiser is allocated more than the desired number of click throughs and no keyword is allocated more than the number of times it appears □

We now show that the allocation has optimal expected welfare.

5.3. *The induced allocation has optimal expected welfare.*

Proof. Suppose for a contradiction that there exists an allocation $A'$ with higher expected welfare. Then $A'$ defines a flow as follows:

- If advertiser $i$ is allocated $u(i)$ units in $A'$ then there is a flow of $u(i)$ on the edge $(s, i)$.
- If advertiser $i$ is allocated in $A'$ $u_j(i)$ impression of advertiser $j$ then there is a flow of $\frac{u_j(i)}{p_{ij}}$ leaving $i$ and going to $j$.
- The flow on $(j, t)$ is defined by $\sum_i p_{ij} u_j(i)$.

It is clear that this is a flow. This flow has cost which is equal to the negated welfare of $A'$ which by definition is higher than the negated welfare of $A$ in contradiction the the definition of $A$. □

These two lemmas provide the desired theorem.

In order to ensure truthfulness we use the standard VCG prices[7] where advertisers pay by critical value. Since this is a multi-unit auction this ensures truthfulness for both the value and the desired number of units. Note that although this mechanism charges for impressions it can easily be modified to charge for clickthoughs.

We then get our main theorem:

3. *There exists a quasi polynomial time, truthful mechanism which maximizes welfare for the keywords problem with budget constraints.*

## 6. REVENUE MAXIMIZATION

Theorem 3 shows that our proposed mechanism optimizes welfare. However, it has been brought to our attention that some of the search engines wish to maximize revenue (see e.g., [14]) . In this section we show that our mechanism actually increases the revenue over the current method of auctioning keywords.

We will show that the allocation that our mechanism $\mathcal{A}$ discovers is an *equilibria* for the *current mechanism* $\mathcal{A}'$ used to allocate keywords. This will then show that the allocation defined by $\mathcal{A}$ is an allocation supported by an equilibria report by the players for the current mechanism $\mathcal{A}'$. This allows us to prove that payments are higher in $\mathcal{A}$ then in the same allocation in $\mathcal{A}'$[8].

We first note that for risk-neutral players there is no difference between paying for a probability of an impression and paying for an impression. We now need to show that any player $i$ allocated the entire demand desired $n_i$ is happy with the keywords chosen and any player $i'$ not allocated $n_{i'}$ is also happy not to be allocated.

6.1. *There is no player i who wishes to increase his allocation allotted in $\mathcal{A}$ given that the allocation and prices are determined by $\mathcal{A}'$.*

Proof. Suppose that such a player exists. Then there is a keyword $j$ that $i$ wants to receive. There are two cases:

1. $i$ does not receive $n_i$ units: In this case we look at the player $i'$ with the lowest payoff for $j$. Now by definition $i'$ payoff for $j$ is higher than $i$'s payoff for $j$ (as otherwise $\mathcal{A}$ which maximizes welfare would have allocated it to $i$) and hence the price that $i$ would pay for $j$ is higher than $i$ value.
2. $i$ receives $n_i$ units: In this case $i$ has no marginal benefit from an increased allocation. We show below that $i$ has no benefit from transferring a unit received from one keyword to another.

□

6.2. *For any player i and any allocation that $\mathcal{A}$ outputs, i will not attempt to reduce his allocation in $\mathcal{A}'$.*

Proof. By the incentive compatibility of $\mathcal{A}$ any $i$ has a positive marginal utility from an unit allocated to $i$. It remains to show that $i$ still has a positive marginal utility in the current mechanism (with the different prices). Since $i$ has a positive value for every keyword $j$ and since (by the previous lemma) for any advertiser $i'$, $i'$ does not want to increase his allocation of keyword $j$ in $\mathcal{A}$, then the payoff for $i'$ must be less than the payoff for $i$ and hence the payment of $i$ must be less than $i$'s utility. □

6.3. *Given an allocation A output by $\mathcal{A}$, there is no player i who prefers to win in different keywords in the current mechanism than the keywords that are allocated to i in A.*

Proof. At first glance this appears obvious inasmuch as $\mathcal{A}$ maximizes welfare and charges VCG prices. However, since the current mechanism charges different prices this might no longer be true.

Suppose for a contradiction that $i$ is an advertiser who wishes to move a unit of allocation from keyword $j$ to keyword $j'$. Denote by $i'$ the loser with the highest payoff in $j$ and by $i''$ the winner with the lowest payoff in $j'$. Since $\mathcal{A}$ optimizes welfare the payoff of $i''$ is more than the payoff of $i'$ and hence in $\mathcal{A}'$ $i$ will pay more when winning $j'$ then when winning $j$. In contradiction. □

Now that we've shown that player will not attempt to change the allocation in $\mathcal{A}'$ if they are allocated based on our mechanism we need to show that there are prices in the current mechanism that can support the allocation induced by our proposed mechanism.

6.4. *For any player i and any allocation A that $\mathcal{A}$ outputs there exists values for players that support A in the $\mathcal{A}'$*

Proof. Since $\mathcal{A}'$ allows different values for different keywords simply set the price to be the critical value for each keyword. □

[7]Since we use augmenting flows, the standard VCG prices can be characterized as distances in the residual flow graph.
[8]Of course, this assumes that the expectation (on number of keywords) is known to the advertisers.

Remark 6.1. *This proof utilizes the fact that in $\mathcal{A}'$ different values can be set for different keywords. It is possible although more difficult to show that the lemma holds even when players are constrained to give a single value for any clickthrough.*

Combining the above lemmas we get the following theorem:

4. *The allocation $A$ output by $\mathcal{A}$ is an equilibria for the $\mathcal{A}'$.*

We now need an additional technical lemma:

6.5. *The price paid by any player $i$ in $\mathcal{A}$ is at least the price paid by $i$ in the equivalent equilibria in the $\mathcal{A}'$.*

Proof. Since the allocation induced by the generalized min cost flow is an equilibria it suffices to show that the prices are higher. Since we have shown that the set of winners (and hence the set of losers) is the same for the induced allocation as well as the equilibria it remains to show that prices are higher.

This follows from the observation that the prices in the current mechanism for any allocation of any keyword is set to be the critical value for winning that keyword. In $\mathcal{A}$ it is set to be the critical value for winning *any* keyword. Since the set of losers is the same obviously the latter value is higher. □

This then yields the desired corollary:

6.6. *There exists an equilibria for $\mathcal{A}'$ whose revenue is no more than the revenue of $\mathcal{A}$.*

Remark 6.2. *The ratio between the current revenue and the revenue in our mechanism might be $O(m)$. It is easy to show by a greedy argument that the ratio can be bounded by $m$. The following example shows that this ratio can indeed be as bad as $O(m)$.*

2. *Let there be $m+1$ players. For all $\forall i : v_i = b_i = 1$For $1 \leq i \leq m$ $p_{ii} = 1-\epsilon$ and otherwise $p_{ij} = 0$. $\forall j : p_{m+1,j} = 1$. $\mathcal{A}'$ will allocate to player $m+1$ as well as to $m-1$ other players. However, the first allocation will be to player $m+1$ which will pay $1-\epsilon$. After player $m+1$ is satisfied there is no competition for any other keyword so the revenue for all keywords beyond the first one is zero.*

*In contrast $\mathcal{A}$ will have revenue of $m(1-\epsilon)$ since the critical value for any player is $1-\epsilon$ (as otherwise a higher flow can be achieved by not allocating to that player).*

# 7. WHEN ADVERTISERS PUT MONETARY BOUNDS ON BUDGETS

In the discussion above we assumed that the advertisers bound the number of click-throughs that they want. In practice, advertisers today bound the budget that they wish to spend on advertising in monetary form. In this section we show that we can translate the budget constraints into a constraint on the number of click-throughs. We show that this transformation can be done in quasi polynomial time.

Our reduction will proceed in stages. We will first start out with each advertiser demanding $\sum_j a_j$ click throughs. Therefore the initial total demand is $n \sum_j a_j$. We will iteratively reduce the demand until no advertiser is over budget. Obviously, if we succeed in reducing the demand and at each stage this yields a quasi-polynomial reduction. Denote by $b_i$ the budget of advertiser $i$.

At each stage $t$:

1. Look at the set $S_t$ of advertisers who currently pay more than $b_i$.

2. For each $i \in S_t$ we independently reduce the demand so that $i$ does not pay more than $b_i$. Obviously, this can be done in quasi-polynomial time.

3. Find the $i'_t \notin S_t$ s.t. the welfare is maximized if we increase the allocation of $i'_t$ by $\min_{i,j} p_{ij}$ and this increase will not cause $i'$ to be over budget.

4. This click through will be along some combination of keywords.

5. Reduce the capacity of the keywords by what is allocated to $i'_t$.

6. Temporarily increase the demand of all $i \in S_t$ to $\sum_j a_j$ and calculate the optimal allocation using $\mathcal{A}$ when constraints are on demands.

7. Set the demand for all $i \in S_t$ to be the amount allocated to $i$ in the previous step.

We first note that since $S_t$ is monotonically decreasing and since the capacity allocated to $S_t$ strictly decreases at any time the algorithm converges. Since the amount of capacity allocated to $S_t$ decrease by $\min_{i,j} p_{ij}$ the algorithm converges in quasi-polynomial time.

Since $i'$ is chosen to be the optimal increase and since if $i'_t$ does not receive this increase then for any $t' > t$ $i'_{t'}$ will not and hence for all $i \in S_t$ the payment will be more than the budget this maintains the optimal welfare property.

Although [5] shows an impossibility result on optimizing welfare in the presence of budgets we circumvent this by assuming that the welfare is optimized when the budget is exhausted (i.e., achieving the optimal welfare achievable by mechanisms when budgets exist).

# 8. CONCLUSIONS AND FUTURE WORK

We showed that current mechanism used by search engines for allocating keywords suffer from a possibly large inefficiency. We constructed a mechanism that rectifies this problem. There are several important practical open questions.

The first question is whether advertisers will agree to participate in a mechanism in which they lose control over which exact keywords they receive in return for a guarantee of higher welfare. Fortunately, we do not require that all of the advertisers agree to participate. We can easily simulate any advertiser who is unwilling to lose control over advertisement placement by setting the relevant probabilities to be zero[9].

[9]It is even possible for advertisers to choose among which keywords they are willing to relinquish control. It is also possible for search engines to "bundle" groups of keywords such that advertisers have no choice what keywords in a given bundle they receive.

A side effect of the ability of advertisers to opt in on a per advertiser basis into this mechanism is that even given current mechanisms for keyword auctions advertisers can minimize their payments while optimizing their value by calculating the equilibria given by our mechanism. This improves current results of optimization on the advertiser level [19, 9]. However, we do assume that advertisers have knowledge of the other advertisers demands as well as the supply of available keywords and the click through rates.

The loss of welfare of the current methods of allocating keywords can be be a factor 2 in theory. In practice it would be interesting to see what the actual loss is given the known click through rates for different advertisers and keywords. It is possible (although unlikely) that the problem that this paper points out which depends on variances in probabilities does not arise in practice.

Our mechanism is quasi-polynomial but finds the optimal allocation, it is also to construct a fully polynomial truthful approximation based on results that appear in [20]. This however requires much more careful analysis. It is unclear how important the question of quasi-polynomial is in practice given the click-through probabilities and the actual running time of the mechanism on real data[10].

[17] pose the intriguing question of whether it is possible to create a mechanism that on average has a revenue of $1 - o(1)$ of the optimal and in the worst case has $\frac{e-1}{e}$ of the optimal. It is tempting to do the following to answer their open question:

1. Set a parameter $p = o(1)$. Run $\mathcal{A}$.

2. If during the running of $\mathcal{A}$ the distribution is unlikely to have come from the original distribution (i.e., there is a probability larger than $p$ that the distributions differ) as defined by some predetermined statistical test then:

3. Run the mechanism defined in [17].

The reason that such a scheme seems workable is that if the distribution is the expected distribution w.h.p. we will continue to think that the distribution is correct while with some small probability we will run [17] at a loss of $\frac{1}{e}$ of the welfare. However, if the distribution is not the expected distribution we will notice this relatively quickly and hence move to [17] to achieve almost an $\frac{e-1}{e}$ approximation (with some loss due to the time required to discover that we are not in the expected distribution).

The problem is that the mechanism defined in [17] is not efficient. This means that the $\frac{e-1}{e}$ approximation of the revenue achieved by [17] might be *higher* than the optimal welfare achieved by our mechanism.

There are two solutions to this problem. The first is to look at a hybrid mechanism which is no longer truthful. The second is to to define a variant of [17] that is $\delta$-*gain truthful* a concept defined recently by [13]. In this variant the equilibrium welfare can indeed be compared to the welfare achieved by our mechanism and therefore we can create such a hybrid mechanism. We are currently exploring this avenue.

## 9. ACKNOWLEDGEMENTS

The author wishes to acknowledge Eva Tardos for helpful remarks and comments. The author wishes to thank Yonatan Amit for helping improve the readability. The author also wishes to thank the reviewers for helpful comments.

## 10. REFERENCES

[1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC 06*.

[2] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Impersonation based mechansisms. In *AAAI 06*.

[3] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and implementation in undominated strategies. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1054–1063, New York, NY, USA, 2006. ACM Press.

[4] Dirk Bergemann and Juuso Vlimki. Efficient dynamic auctions. Technical report, October 2006.

[5] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 44–51, New York, NY, USA, 2005. ACM Press.

[6] Peter Cramto, Yoav Shoham, and Richard Steinberg (editors). Combinatorial auctions. In *MIT Press*.

[7] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 610–618, New York, NY, USA, 2005. ACM Press.

[8] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In *NBER Working Papers 11765, National Bureau of Economic Research, Inc.*

[9] Jon Feldman, S. Muthukrishnan, Martin Pal, and Cliff Stein. Budget optimization in search-based advertising auctions. In *Manuscript*, 2006.

[10] R. Gonen and D. Lehmann. Linear programming helps solving large multi-unit combinatorial auctions, 2001.

[11] Rica Gonen and Daniel J. Lehmann. Optimal solutions for multi-unit combinatorial auctions: branch and bound heuristics. In *ACM Conference on Electronic Commerce*, pages 13–20, 2000.

[12] Rica Gonen and Elan Pavlov. An incentive compatible multi armed bandit mechanism. In *Submitted for publication*.

[13] Rica Gonen and Elan Pavlov. A truthful adaptive algorithm for allocating keywords with time and budget constraints. In *Manuscript*, 2007.

[14] Google Inc. Sec file 333-140498, 2007.

[15] Brendan Kitts and Benjamin Leblanc. Optimal bidding on keyword auctions. In *Electronic Markets Volum 14(3) pp 186-201*.

[16] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, 2002.

[17] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized on-line matching. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 264–273, Washington, DC, USA, 2005. IEEE Computer Society.

[18] Amir Ronen and Daniel Lehmann. Nearly optimal multi attribute auctions. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 279–285, New York, NY, USA, 2005. ACM Press.

[19] Paat Rusmevichientong and David P. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 260–269, New York, NY, USA, 2006. ACM Press.

[20] Kevin Wayne. A polynomial combinatorial algorithm for generalized minimum cost flow. In *Mathematics of operations research Vol. 27, No. 3, August 2002, pp. 445-459.*

[10] There are many cases in which the bounds given on quasi-polynomial algorithms are much weaker than the actual performance.