

RANDOMIZED METHODS FOR SOLVING THE WINNER DETERMINATION PROBLEM IN COMBINATORIAL AUCTIONS

Joshua C. C. Chan

Department of Mathematics
University of Queensland
Brisbane, QLD 4072, AUSTRALIA

Dirk P. Kroese

Department of Mathematics
University of Queensland
Brisbane, QLD 4072, AUSTRALIA

ABSTRACT

Combinatorial auctions, where buyers can bid on bundles of items rather than bidding them sequentially, often lead to more economically efficient allocations of financial resources. However, the problem of determining the winners once the bids are submitted, the so-called Winner Determination Problem (WDP), is known to be NP hard. We present two randomized algorithms to solve this combinatorial optimization problem. The first is based on the Cross-Entropy (CE) method, a versatile adaptive algorithm that has been successfully applied to solve various well-known difficult combinatorial optimization problems. The other is a new adaptive simulation approach by Botev and Kroese, which evolved from the CE method and combines the adaptiveness and level-crossing ideas of CE with Markov Chain Monte Carlo techniques. The performance of the proposed algorithms are illustrated by various examples.

1 INTRODUCTION

An auction is a common method of determining the value of commodities that have an undetermined price. This mechanism is often used when there is a large number of bidders who are interested in acquiring certain items. These items are auctioned in sequence until they are all sold. In this type of auction, determination of the winner is trivial, as the highest bid “wins” (gets the item). However, in many situations a bidder’s valuation for a combination of items for sale is not the sum of the individual items’ valuations. It is because related assets often have what economists call complementarity and substitution effects, and economic efficiency is improved if bidders are allowed to bid on *combinations* of different assets, rather than bidding them sequentially. Complementarity occurs when two items have features that complement each other, and their combined value for the bidder is higher than the sum of the individual values. Substitutability is the opposite situation, where two items substitute each other, and their combined value is

less than the sum of individuals’. To avoid the problem of complementarity and substitutability, *combinatorial auction* is introduced where buyers bid on sets of items, instead of single items.

Combinatorial auction was first proposed by Jackson (1976) for the allocation of radio spectrum. Rassenti et al. (1982) applied the idea to auctioning airport time slots. Other examples include Federal Communications Commission auctions for wireless communication spectrum (Cramton 1998), auctions for railroad segments (Brewer 1999), and applications in electronic business (Narahari and Dayama 2005). A recent survey on the topic is given by de Vries and Vohra (2003); Cramton et al. (2006) wrote a comprehensive book on the topic that also deals with other aspects of combinatorial auctions. Despite its potentially wide applicability, combinatorial auction presents some new challenges. The most obvious one being the so-called *Winner Determination Problem* (WDP): how to efficiently determine the winner once the bids have been submitted to the auctioneer? Since each bid in a combinatorial auction can be made on an arbitrary set of items, the number of bids grow exponentially with the number of items, making the evaluation difficult for a relatively large auction. In fact, it is known that the problem is NP-hard (de Vries and Vohra 2003).

The purpose of this paper is to present two different randomized algorithms to solve two versions of the WDP. The first is based on the Cross-Entropy (CE) method, a versatile adaptive algorithm that has been successfully applied to solve various well-known difficult combinatorial optimization problems. The other is an adaptive simulation approach proposed by Botev and Kroese (2008), referred to here as the BK method, that evolved from the CE method. We demonstrate the performance of the proposed algorithms by various examples. The rest of the article is organized as follows. We formally introduce the WDP in Section 2 and discuss the CE and the BK methods in Sections 3 and 4 respectively. Numerical experiments are presented in Section 5 and the last section concludes the article.

2 THE WINNER DETERMINATION PROBLEM

The central problem arising from combinatorial auctions is winner determination, which is described as follows. Suppose an auctioneer has a collection of items to auction to a number of bidders, who submit bids on every combination of items. Given the set of bids, the auctioneer then determines the allocation of items to bidders that maximizes her revenue under the constraint that each item is allocated to at most one bidder (with the possibility that the item is not sold). This problem can be formally stated as a combinatorial optimization problem in the following way. Let $I = \{1, \dots, m\}$ be the set of distinct items to be auctioned and $J = \{1, \dots, n\}$ the set of bidders. For every subset $S \subseteq I$, each bidder $j \in J$ is to submit a bid, denoted as $b^j(S)$, he is willing to pay. Let $x(S, j)$ be equal to 1 if the bundle $S \subseteq I$ is allocated to $j \in J$, and $x(S, j) = 0$ otherwise. The WDP can be formulated as:

$$\max \sum_{j \in J} \sum_{S \subseteq I} b^j(S) x(S, j) \quad (1)$$

subject to

$$\begin{aligned} \sum_{S: i \in S} \sum_{j \in J} x(S, j) &\leq 1, \quad \forall i \in I, \\ \sum_{S \subseteq I} x(S, j) &\leq 1, \quad \forall j \in J, \\ x(S, j) &\in \{0, 1\}, \quad \forall S \subseteq I, j \in J. \end{aligned}$$

The first constraint ensures that each object is allocated at most once while the second constraint ensures that no bidder receives more than one bundle. For convenience, let \mathcal{E}_1 be the set of all feasible allocations. Thus the problem becomes maximizing $H_1(x) = \sum_{j \in J} \sum_{S \subseteq I} b^j(S) x(S, j)$ over the set \mathcal{E}_1 . Call this problem WDP1. The most obvious difficulty of such an auction is that bidders are required to submit a bid for every subset of items he is interested in. Thus for an auction with m items a bidder is required to submit $2^m - 1$ bids. Nisan (2000) discusses various ways in which bids can be restricted and their consequence. Even when this enumeration problem can be resolved satisfactorily, WDP1 is still a non-trivial problem. In fact, the above formulation is an instance of the Set Packing Problem (SPP), a well-studied integer programming problem that is known to be NP-hard (de Vries and Vohra 2003).

We present a simple example for illustration. Suppose an auctioneer has $m = 3$ items ($\{b_1, b_2, b_3\}$) to be auctioned, and there are $n = 4$ bidders ($\{A_1, A_2, A_3, A_4\}$), whose bids are given in Table 1. One possible solution is to allocate items $\{b_1, b_3\}$ to bidder A_1 and item $\{b_2\}$ to bidder A_4 , which earns a revenue of $2 + 3 = 5$. For this simple example there are altogether $3^4 = 81$ feasible allocations. In general, for a combinatorial auction with m items and n bidders,

Table 1: Bids submitted for WDP1 with $n = 4$ bidders and $m = 3$ items to be auctioned.

Subset	A_1	A_2	A_3	A_4
$\{b_1\}$	1	2	3	2
$\{b_2\}$	1	1	2	3
$\{b_3\}$	2	1	1	2
$\{b_1, b_2\}$	3	4	1	3
$\{b_1, b_3\}$	2	4	6	4
$\{b_2, b_3\}$	5	3	5	3
$\{b_1, b_2, b_3\}$	6	7	4	5

there are altogether n^m allocations, making the problem intractable when either m or n is large.

Since in many empirical applications the number of items to be auctioned is large and putting a price on every subset is simply impractical, some researchers looked at another type of combinatorial auction described as follows. Suppose the auctioneer has a set of distinct items, $I = \{1, \dots, m\}$, to sell, and each buyer, instead of submitting a bid on every subset of I , submits only one bid, denoted as B_j , $j = 1, \dots, n$. A bid is a tuple $B_j = (S_j, p_j)$, where $S_j \subseteq I$ is a combination of items the buyer is interested in and p_j , $p_j \geq 0$ is a price. Given the bids B_1, \dots, B_n , the auctioneer then decide which bids as winning so as to maximize her revenue under the constraint that each item can be allocated to at most one bidder. Formally, the problem can be formulated as

$$\max \sum_{j \in J} p_j x_j \quad (2)$$

subject to

$$\begin{aligned} \sum_{j: i \in S_j} x_j &\leq 1, \quad \forall i \in I, \\ x_j &\in \{0, 1\}. \end{aligned}$$

This problem is also known to be NP hard. For convenience, let \mathcal{E}_2 be the set of all feasible allocations. Then the problem becomes maximizing $H_2(x) = \sum_{j \in J} p_j x_j$ over the set \mathcal{E}_2 . Call this problem WDP2. Sandholm et al. (2005) provides a fast algorithm to solve this version of WDP.

3 COMBINATORIAL OPTIMIZATION VIA CE

In this section we discuss a CE based algorithm to solve the combinatorial optimization problems in (1) and (2). We follow the approach proposed in Rubinstein and Kroese (2004). Suppose we wish to maximize a function $H(\mathbf{x})$ over some discrete set \mathcal{E} . Denote the maximum by γ^* , that is,

$$\gamma^* = \max_{\mathbf{x} \in \mathcal{E}} H(\mathbf{x}). \quad (3)$$

Instead of seeking a solution to the optimization problem in (3), the CE method recasts the original problem into an estimation problem of rare-event probabilities. By doing so, it aims to obtain an optimal parametric sampling distribution on \mathcal{E} , rather than finding the optimal solution directly. To this end, let $\{I_{\{H(\mathbf{x}) \geq \gamma\}}\}$ be a collection of indicator functions on \mathcal{E} for various levels $\gamma \in \mathbb{R}$. Next, let $\{f(\cdot; \mathbf{v}), \mathbf{v} \in \mathcal{V}\}$ be a family of probability density functions on \mathcal{E} parametrized by a real-valued parameter vector \mathbf{v} . For a fixed vector $\mathbf{u} \in \mathcal{V}$ we associate with (3) the problem of estimating the rare-event probability

$$\ell(\gamma) = \mathbb{P}_{\mathbf{u}}(H(\mathbf{X}) \geq \gamma) = \mathbb{E}_{\mathbf{u}}[I_{\{H(\mathbf{X}) \geq \gamma\}}], \quad (4)$$

where $\mathbb{P}_{\mathbf{u}}$ is the probability measure under which the random state \mathbf{X} has a discrete pdf $f(\cdot; \mathbf{u})$ and $\mathbb{E}_{\mathbf{u}}$ denotes the corresponding expectation operator. Then the goal of the CE method is to generate a sequence of pdfs $f(\cdot; \hat{\mathbf{v}}_0), f(\cdot; \hat{\mathbf{v}}_1), \dots$ converging to a degenerate measure (Dirac measure) that assigns all probability mass to a single state \mathbf{x}_T , for which, ideally, the function value is the global optimal.

More specifically, we start with a parametrized sampling distribution $f(\cdot; \hat{\mathbf{v}}_0)$ from which a random sample of size N , $\mathbf{X}_1, \dots, \mathbf{X}_N$, is generated. For each observation \mathbf{X}_i , we compute its performance, denoted as $H(\mathbf{X}_i)$. A fixed number of the best observations are singled out and referred to as the *elite sample*. Operationally, we first compute the sample $(1 - \rho)$ -quantile of the performances as

$$\hat{\gamma}_t = H_{(\lceil (1-\rho)N \rceil)}, \quad (5)$$

where $H_{(i)}$ is the i -th ordered observation and $\lceil \cdot \rceil$ is the ceiling function. Then we select all the observations \mathbf{X}_i for which $\mathbf{X}_i \geq \hat{\gamma}_t$ and call this collection the elite sample. The elite sample is then used to update the parameters for the sampling distribution by solving the following maximization problem

$$\max_{\mathbf{v}} D(\mathbf{v}) := \max_{\mathbf{v}} \frac{1}{N} \sum_{i=1}^N I_{\{H(\mathbf{X}_i) \geq \hat{\gamma}_t\}} \log f(\mathbf{X}_i; \mathbf{v}). \quad (6)$$

Instead of updating the parameter vector \mathbf{v} directly via the solution of (6), it is often better to use the following smoothed version

$$\hat{\mathbf{v}}_t = \alpha \tilde{\mathbf{v}}_t + (1 - \alpha) \hat{\mathbf{v}}_{t-1}, \quad (7)$$

where $\tilde{\mathbf{v}}_t$ is the parameter vector obtained from the solution of the maximization program (6) and $\alpha \in (0, 1)$ is referred to as the *smoothing parameter*. Thus the CE algorithm for optimization can be summarized as follows.

Algorithm 1 [CE Algorithm for Optimization]

1. Choose an initial parameter vector $\hat{\mathbf{v}}_0$. Set $t = 1$.
2. Generate a sample $\mathbf{X}_1, \dots, \mathbf{X}_N$ from the density $f(\cdot; \hat{\mathbf{v}}_{t-1})$ and compute the sample $(1 - \rho)$ -quantile $\hat{\gamma}_t$ of the performance according to (5).
3. Use the same sample $\mathbf{X}_1, \dots, \mathbf{X}_N$ to solve the stochastic program (6) and denote the solution as $\tilde{\mathbf{v}}_t$.
4. Apply (7) to smooth out the vector $\tilde{\mathbf{v}}_t$.
5. If for some $t = d$, say $d = 4$,

$$\hat{\gamma}_t = \hat{\gamma}_{t-1} = \dots = \hat{\gamma}_{t-d},$$

then stop; otherwise set $t = t + 1$ and iterate from Step 2.

To implement the above CE program to solve the combinatorial optimization problems (1) and (2), we first need a convenient way to represent the set of all feasible allocations \mathcal{E}_1 and \mathcal{E}_2 and assign a family of pdfs on them. For WDP1 with n buyers and m items to be auctioned, a feasible allocation can be uniquely represented by a vector $\mathbf{X} = (X_1, \dots, X_m)$, where $X_i = j$ indicates that item i is allocated to bidder j . For example, for a combinatorial auction with 6 items and 10 bidders ($m = 6$ and $n = 10$), the vector $(3, 9, 8, 9, 3, 3)$ represents an allocation of items $\{1, 5, 6\}$ to bidder 3, items $\{2, 4\}$ to bidder 9 and item $\{3\}$ to bidder 8. Given this convenient representation, a random allocation can be easily generated from the set \mathcal{E}_1 simply by letting X_1, \dots, X_m be independent n -point random variables such that $\mathbb{P}(X_i = j) = p_{ij}$, $i = 1, \dots, m$, $j = 1, \dots, n$. In addition, the stochastic program in (6) can be solved analytically. Specifically, given the sample $\mathbf{X}_1, \dots, \mathbf{X}_N$ we update the components of $\hat{P}_t = (\hat{p}_{t,ij})$ as

$$\hat{p}_{t,ij} = \frac{\sum_{k=1}^N I_{\{H(\mathbf{X}_k) \geq \hat{\gamma}_t\}} I_{\{X_{ki}=j\}}}{\sum_{k=1}^N I_{\{H(\mathbf{X}_k) \geq \hat{\gamma}_t\}}}.$$

For WDP2 with n bids, an allocation can be represented by a vector $\mathbf{X} = (X_1, \dots, X_n)$, where $X_i = 1$ if buyer i wins the bid and 0 otherwise. For example, for a combinatorial auction with $n = 5$ bids, the vector $\mathbf{X} = (1, 0, 0, 1, 1)$ represents the allocation that buyers 1, 4 and 5 win their bids while buyers 2 and 3 do not. It is worth noting that some of these vectors might represent infeasible allocations. For instance, if bids 1 and 2 both contain the same item, then $(1, 1, \dots)$ is not a feasible allocation as one item cannot be allocated to both buyers. In that case we can simply set the performance of such an allocation to be 0. Given this representation, a random allocation can be easily generated simply by letting each component of $\mathbf{X} = (X_1, \dots, X_n)$ be an independent Bernoulli random variable such that the success probability of X_i is p_i . Finally, given the sample $\mathbf{X}_1, \dots, \mathbf{X}_N$, the updating rule

is simply

$$\hat{p}_{t,i} = \frac{\sum_{k=1}^N I_{\{H(\mathbf{X}_k) \geq \hat{\gamma}\}} I_{\{X_{ki}=1\}}}{\sum_{k=1}^N I_{\{H(\mathbf{X}_k) \geq \hat{\gamma}\}}}.$$

4 COMBINATORIAL OPTIMIZATION VIA BK

In this section we discuss an adaptive simulation approach proposed by [Botev and Kroese \(2008\)](#), referred to as the BK method, that has proved to be very useful for rare-event probability estimation, combinatorial optimization and counting problems. Evolved from the CE method, the BK method circumvents the likelihood ratio degeneracy problem by sampling directly from the minimum-variance importance density. Specifically, the BK method first associates the optimization problem (3) with the corresponding rare-event probability estimation problem (4). But instead of constructing a sequence of pdfs $f(\cdot; \hat{\gamma}_0), f(\cdot; \hat{\gamma}_1), \dots$ converging to a degenerate measure as in CE, the BK method aims to directly sample from the minimum-variance importance sampling density

$$g^*(\mathbf{x}|\gamma) = \frac{f(\mathbf{x}; \mathbf{u}) I_{\{H(\mathbf{x}) \geq \gamma\}}}{\ell(\gamma)}. \quad (8)$$

As opposed to estimation problems, we are not interested in obtaining an unbiased estimate for the rare-event probability *per se*. Rather, we only wish to approximately sample from the pdf (8) for as large a value of γ as possible.

Suppose we have a sequence of levels $-\infty = \gamma_0 < \gamma_1 < \dots < \gamma_T = \gamma^*$ judiciously chosen. Heuristically, if we generate N draws from $\mathbf{X}_i^{(0)} \sim g^*(\mathbf{x}|\gamma_0) = f(\mathbf{x}; \mathbf{u})$ and accept only those $\mathbf{X}_i^{(0)}$ for which $\mathbf{X}_i^{(0)} \geq \gamma_1$, then we have a sample from $g^*(\mathbf{x}|\gamma_1)$. Call this sample $\mathcal{X}^{(1)}$ and let N_1 be the sample size. Presumably $N_1 < N$. To obtain N draws from $g^*(\mathbf{x}|\gamma_1)$, we first sample uniformly with replacement N times from the population $\mathcal{X}^{(1)}$ to acquire a new sample $\mathbf{Y}_1, \dots, \mathbf{Y}_N$. For each $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{ik})$, we then sample each component Y_{ij} from $g^*(\mathbf{x}|\gamma_1)$ *conditionally* on the other components. Since the resampling and conditional sampling steps do not change the underlying distribution, in this way, we obtain a sample $\mathbf{X}_1^{(1)}, \dots, \mathbf{X}_N^{(1)}$ from $g^*(\mathbf{x}|\gamma_1)$. Now let $\mathcal{X}^{(2)}$ denote the collection of $\mathbf{X}_i^{(1)}$ such that $\mathbf{X}_i^{(1)} \geq \gamma_2$. Obviously $\mathcal{X}^{(2)}$ is a sample from $g^*(\mathbf{x}|\gamma_2)$. We then apply the resampling and conditional sampling procedures to obtain a new sample of size N from $g^*(\mathbf{x}|\gamma_2)$. We repeat this process until we reach γ^* . The BK algorithm for optimization is summarized below.

Algorithm 2 [BK Algorithm for Optimization]

1. Set $t = 1$. Generate a sample $\mathcal{X}^{(0)} = \{\mathbf{X}_1^{(0)}, \dots, \mathbf{X}_N^{(0)}\}$ from the nominal density $f(\mathbf{x}; \mathbf{u})$ and compute the sample $(1 - \rho)$ -quantile $\hat{\gamma}_0$ of the performance according to (5). Let $\tilde{\mathcal{X}}^{(0)} = \{\tilde{\mathbf{X}}_1^{(0)}, \dots, \tilde{\mathbf{X}}_{N_0}^{(0)}\}$ be the subset of the population $\{\mathbf{X}_1^{(0)}, \dots, \mathbf{X}_N^{(0)}\}$ for which $H(\mathbf{X}_i^{(0)}) \geq \hat{\gamma}_0$. Then we have

$$\tilde{\mathbf{X}}_1^{(0)}, \dots, \tilde{\mathbf{X}}_{N_0}^{(0)} \stackrel{appr.}{\sim} g^*(\mathbf{x}|\hat{\gamma}_0).$$

2. Sample uniformly with replacement N times from the population $\tilde{\mathcal{X}}^{(t-1)}$ to obtain a new sample $\mathbf{Y}_1, \dots, \mathbf{Y}_N$.
3. For each $\mathbf{Y} = (Y_1, \dots, Y_k)$ in $\{\mathbf{Y}_1, \dots, \mathbf{Y}_N\}$, generate $\tilde{\mathbf{Y}} = (\tilde{Y}_1, \dots, \tilde{Y}_k)$ as follows:
 - (a) Draw \tilde{Y}_1 from the conditional density $g^*(y_1|\hat{\gamma}_{t-1}, Y_2, \dots, Y_k)$.
 - (b) Draw $\tilde{Y}_i \sim g^*(y_i|\hat{\gamma}_{t-1}, \tilde{Y}_1, \dots, \tilde{Y}_{i-1}, Y_{i+1}, \dots, Y_k)$, $i = 2, \dots, k-1$.
 - (c) Draw $\tilde{Y}_k \sim g^*(y_k|\hat{\gamma}_{t-1}, \tilde{Y}_1, \dots, \tilde{Y}_{k-1})$.

Denote the resulting population of $\tilde{\mathbf{Y}}$ by $\tilde{\mathbf{X}}_1^{(t)}, \dots, \tilde{\mathbf{X}}_{N_t}^{(t)}$.

4. Let $\hat{\gamma}_t$ be the $(1 - \rho)$ sample quantile of $H(\mathbf{X}_1^{(t)}), \dots, H(\mathbf{X}_{N_t}^{(t)})$ and $\tilde{\mathcal{X}}^{(t)} = \{\tilde{\mathbf{X}}_1^{(t)}, \dots, \tilde{\mathbf{X}}_{N_t}^{(t)}\}$ be the subset of the population $\{\mathbf{X}_1^{(t)}, \dots, \mathbf{X}_{N_t}^{(t)}\}$ for which $H(\mathbf{X}_i^{(t)}) \geq \hat{\gamma}_t$. Again, we have

$$\tilde{\mathbf{X}}_1^{(t)}, \dots, \tilde{\mathbf{X}}_{N_t}^{(t)} \stackrel{appr.}{\sim} g^*(\mathbf{x}|\hat{\gamma}_t).$$

5. If for some $t = d$, say $d = 4$,

$$\hat{\gamma}_t = \hat{\gamma}_{t-1} = \dots = \hat{\gamma}_{t-d},$$

then stop; otherwise set $t = t + 1$ and iterate from Step 2.

6. Deliver the vector \mathbf{X}^* from the set $\mathbf{X}_1^{(T)}, \dots, \mathbf{X}_{N_T}^{(T)}$ for which $H(\mathbf{X}_i^{(T)})$ is maximal as an estimate for the global maximizer of (3).

To implement the above algorithm, we need a convenient way to represent the set of all feasible allocations \mathcal{E}_1 and \mathcal{E}_2 and assign a nominal pdf on each of them. Here we use the same representation as in Section 3: for WDP1 an allocation in \mathcal{E}_1 can be represented by a vector $\mathbf{X} = (X_1, \dots, X_m)$, where $X_i = j$ indicates that item i is allocated to bidder j ; for WDP2 a typical element in \mathcal{E}_2 is a vector $\mathbf{X} = (X_1, \dots, X_n)$, where $X_i = 1$ if buyer i wins the bid and 0 otherwise.

Sampling from the conditional densities in Step 3 is also straightforward. For WDP1, sampling a random variable \tilde{Y}_i

Table 2: Synthetic WDP1 example with $m = 10$, $n = 50$, and $\rho = 0.1$.

method	γ^*	mean	min	max	CPU	\bar{T}
CE	19.58	19.58	19.55	19.58	0.8	13.6
BK	19.58	19.58	19.58	19.58	0.8	20.6

Table 3: Synthetic WDP1 example with $m = 15$, $n = 100$ and $\rho = 0.1$.

method	γ^*	mean	min	max	CPU	\bar{T}
CE	29.75	29.75	29.72	29.75	2.6	18.7
BK	29.75	29.75	29.75	29.75	5.2	37.1

from the conditional density can be accomplished as follows. Let Z be a random variable that follows an n -point distribution. If $H([y_1, \dots, y_{i-1}, Z, y_{i+1}, \dots, y_m]) \geq \hat{\gamma}_{i-1}$, then set $\tilde{Y}_i = Z$; otherwise set $\tilde{Y}_i = Y_i$. For WDP2, the step is analogous: Let $Z \sim \text{Ber}(0.5)$. If $H([y_1, \dots, y_{i-1}, Z, y_{i+1}, \dots, y_m]) \geq \hat{\gamma}_{i-1}$, then set $\tilde{Y}_i = Z$; otherwise set $\tilde{Y}_i = 1 - Z$.

5 NUMERICAL RESULTS

To demonstrate the performance of the proposed CE and BK algorithms, we present various numerical examples in this section. For the WDP1 in (1), we consider two synthetic examples, one with $n = 50$ bidders and $m = 10$ items to be auctioned and the other with $n = 100$ and $m = 15$. Recall that the total number of possible allocations is n^m (i.e., 9.8×10^{16} for the first example and 10^{30} for the second). In the synthetic examples, the bid for each bundle is generated randomly from a uniform distribution $U[10(k-1), 10(k+1)]$, where k is the number of items in the bundle. For the CE method, we set the sample size to be $N = 5,000$, the rarity parameter $\rho = 0.1$ and smoothing parameter $\alpha = 0.9$. For the BK method, we set $N = 500$ and $\rho = 0.1$. The experiments were executed using Matlab 7.4 on a desktop with a 2.66GHz Intel Core 2 Duo CPU. Each algorithm was run 10 times and we report the true global maximum γ^* , together with the mean, minimum and maximum estimated $\hat{\gamma}^*$ in Tables 2 and 3. We also report the mean number of iterations \bar{T} and CPU times (in seconds) needed for each run.

To monitor the convergence property of the CE algorithm, we introduce the quantity

$$P_t^{mm} = \min_{1 \leq i \leq m} \max_{1 \leq j \leq n} \hat{\rho}_{t,ij}.$$

A typical evolution of the CE algorithm for the WDP1 synthetic example with $m = 10$ and $n = 50$ is reported in Table 4.

For WDP2, we also consider two synthetic examples, one with $n = 50$ bids and $m = 30$ items to be auctioned and

Table 4: Evolution of the CE algorithm for the WDP1 example with $m = 10$, $n = 50$, $N = 5,000$ and $\rho = 0.1$.

t	$\hat{\gamma}_t$	$H_{t,(1)}$	P_t^{mm}
1	12.47	16.64	0.04
2	14.47	17.70	0.06
3	16.15	18.23	0.10
4	17.45	19.03	0.14
5	18.29	19.23	0.19
6	18.87	19.45	0.23
7	19.25	19.52	0.27
8	19.43	19.56	0.43
9	19.50	19.57	0.49
10	19.53	19.58	0.59
11	19.55	19.58	0.91
12	19.58	19.58	0.97
13	19.58	19.58	0.98
14	19.58	19.58	0.99
15	19.58	19.58	0.99

Table 5: Synthetic WDP2 example with $m = 30$, $n = 50$, and $\rho = 0.1$.

method	γ^*	mean	min	max	CPU	\bar{T}
CE	4.15	4.09	4.01	4.15	0.6	7.8
BK	4.15	4.15	4.15	4.15	1.0	8.4

Table 6: Synthetic WDP2 example with $m = 50$, $n = 100$ and $\rho = 0.1$.

method	γ^*	mean	min	max	CPU	\bar{T}
CE	5.49	5.48	5.47	5.49	1.3	8.4
BK	5.49	5.49	5.49	5.49	2.7	8.9

the other with $n = 100$ and $m = 50$. Recall that the total number of possible allocations is 2^n (i.e., 1.13×10^{15} for the first example and 1.26×10^{30} for the second). In the synthetic examples, the bids are generated as follows. First pick the number of items randomly from $\{1, \dots, m\}$. Given the number of items, choose randomly that many items without replacement from $\{1, \dots, m\}$. Lastly, generate the price for the bundle from a uniform distribution $U[0, 1]$. For the CE method, we set the sample size to be $N = 7,000$ for the first example and $N = 10,000$ for the second; for the BK method, we set $N = 500$ for both examples. Each algorithm was run 10 times and we report the true global maximum γ^* , together with the mean, minimum and maximum estimated $\hat{\gamma}^*$ in Tables 5 and 6.

These examples suggest that for both versions of the WDP, the CE and the BK methods are fast, efficient and able to locate the global maximum accurately. In fact, the BK method is able to locate the global maximum in every run. On the other hand, even though the CE method is less

robust, it is much faster and its estimate is still very close to the global maximum when it misses it.

6 CONCLUDING REMARKS

In this article we formulate two versions of the Winner Determination Problem in combinatorial auctions and introduce two different adaptive simulation approaches, namely the CE and the BK methods, to solve both problems. We demonstrate the empirical performance of the proposed algorithms by various synthetic examples and the results suggest that the proposed methods are both efficient and able to locate the maximum rather quickly.

ACKNOWLEDGMENTS

This work is supported by the Australian Research Council (Discovery Grant DP0558957). The first author would also like to acknowledge financial support from the Australian Government and the University of Queensland through IPRS and UQRS scholarships.

REFERENCES

- Botev, Z. I., and D. P. Kroese. 2008. An efficient algorithm for rare-event probability estimation, combinatorial optimization, and counting. *Methodology and Computing in Applied Probability* 10(4). To appear.
- Brewer, P. J. 1999. Decentralized computation procurement and computational robustness in a smart market. *Economic Theory* 13:41–92.
- Cramton, P. 1998. The efficiency of the fcc spectrum auctions. *Journal of Law and Economics* 41:727–736.
- Cramton, P., Y. Shoham, and R. Steinberg. 2006. *Combinatorial auctions*. The MIT Press.
- de Vries, S., and R. Vohra. 2003. Combinatorial auctions: A survey. *INFORMS Journal on Computing* 15(3).
- Jackson, C. 1976. *Technology for spectrum markets*. Ph.D. thesis, Department of Electrical Engineering, School of Engineering, MIT.
- Narahari, Y., and P. Dayama. 2005. Combinatorial auctions for electronic business. *Sadhana Special Issue on Electronic Commerce* 30:179–211.
- Nisan, N. 2000. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, 1–12.
- Rassenti, S. J., V. L. Smith, and R. L. Bulfin. 1982. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics* 13:402–417.
- Rubinstein, R. Y., and D. P. Kroese. 2004. *The cross-entropy method: A unified approach to combinatorial optimization monte-carlo simulation, and machine learning*. New York: Springer-Verlag.
- Sandholm, T., S. Suri, A. Gilpin, and D. Levine. 2005. Cabob: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Science* 51:374–390.

AUTHOR BIOGRAPHIES

Joshua C. C. Chan is a PhD student in the department of Mathematics at the University of Queensland, Australia. He completed his undergraduate degree in Economics at the University of Macau, China, and obtained his Master of Science degree in Mathematics at the University of California, Irvine, USA. His research interests include Bayesian statistics and econometrics, simulation and optimization in economics and finance applications. His email address is <chancc@maths.uq.edu.au>.

Dirk P. Kroese is a pioneer of the well-known Cross-Entropy method and coauthor (with R.Y. Rubinstein) of the first monograph on this method. He is the author of a wide range of publications in applied probability and simulation, including the recent Second Edition of *Simulation and the Monte Carlo Method*, Wiley, 2007 (with R.Y. Rubinstein). He is associate editor of *Methodology and Computing in Applied Probability* and guest editor of *Annals of Operations Research*. He has held research and teaching positions at Princeton University and the University of Melbourne, and is currently working at the Department of Mathematics of the University of Queensland. His email and web addresses are <kroese@maths.uq.edu.au> and <<http://www.maths.uq.edu.au/~kroese/>>.