Autonomous Bidding Agents:
Strategies and Lessons from TAC Travel

Amy Greenwald
Brown University

with

Seong Jae Lee and Victor Naroditskiy (Students)
Mike Wellman and Peter Stone (Coauthors)

April 10, 2008

Autonomous Trading Agents

Background

- Trade is a quintessential human activity
- Compared to other interactive decision-making domains, trade is particularly amenable to automation
- Automonous agents (*a.k.a.* agents) are software programs that make decisions without direct human intervention
  - a program that carries out the direction "bid $99 for eBay item 123" is not autonomous
  - a program that carries out the direction "buy a digital camera at a good price" is autonomous

Trading Agent Research

Aim

- develop techniques for the effective design and analysis of trading agents
  - specific solutions to particular trading problems, as well as
  - general principles to guide the development of trading agents

Challenge

- markets are multiagent environments, in which the performance of a particular agent's strategy depends on the other agents' behavior
- natural approach is for separate institutions/researchers to develop agents to participate in a common market environment

Trading Agent Research

Aim

- develop techniques for the effective design and analysis of trading agents
  - specific solutions to particular trading problems, as well as
  - general principles to guide the development of trading agents

Challenge

- markets are multiagent environments, in which the performance of a particular agent's strategy depends on the other agents' behavior
- natural approach is for separate institutions/researchers to develop agents to participate in a common market environment

Enter: The Trading Agent Competition

Overview

I. TAC Travel Market Game

- Rules
- Anecdotes

II. Agent Architecture

- Prediction
- Optimization

ebay Auctions

Sequential & Simultaneous

Combinatorial Valuations

- Complementary Goods
  - $v(A) + v(B) \leq v(A \cup B)$
  - camera, flash, and tripod
- Substitutable Goods
  - $v(A) + v(B) \geq v(A \cup B)$
  - Canon AE-1 and Canon A-1

The eBay Bidding Problem

Key Question

Trading Agent: "How do I bid in separate markets for goods whose values are highly interdependent?"

Our Solution

Agent Architecture (Prediction + Optimization), studied in the context of the Trading Agent Competition

I. TAC Travel Market Game

An Example

- Simultaneous & Sequential Auctions
- Complementary and Substitutable Goods

Rules

Complementary and Substitutable Goods

- Flights: Inbound and Outbound
- Hotels: Grand Hotel and Le FleaBag Inn
- Entertainment: Red Sox, Symphony, Theatre

Rules

Simultaneous Auctions

- Flights: infinite supply, prices follow random walk, clear continuously, no resale permitted
- Hotels: ascending, multi-unit, 16th price auctions, random auction closes each minute, no resale permitted
- Entertainment: continuous double auctions, initial endowment, resale is permitted

Rules

Feasible Packages

- arrival date prior to departure date
- same hotel on all intermediate nights
- at most one entertainment event per night
- at most one of each type of entertainment

Rules

Client Preferences

Client	IAD	IDD	HV	R	S	T
1	1	3	99	134	118	65
2	1	4	131	170	47	49
3	1	2	147	13	55	49
4	3	4	145	130	60	85
5	1	4	82	136	68	87
6	2	4	53	94	51	105
7	1	3	54	156	126	71
8	1	5	113	119	187	143

Rules

Utility = 1000 - travelPenalty + hotelBonus + funBonus

$$
\begin{aligned}
\text{travelPenalty} &= 100(|\text{IAD} - \text{AD}| + |\text{IDD} - \text{DD}|) \\
\text{hotelBonus} &= \begin{cases} \text{HV} & \text{if H = G} \\ 0 & \text{otherwise} \end{cases} \\
\text{funBonus} &= \text{entertainment values}
\end{aligned}
$$

Anecdotes

- 2000: `eBay` sniping
- 2001: LivingAgents
- 2002: WhiteBear
- 2003: TAC SCM
- 2004: WhiteBear
- 2005: Mertacor
- 2006: RoxyBot

II. Agent Architecture

LOOP

0. Download current prices and winnings from server
1. PREDICT build a model of the auctions' clearing prices
2. OPTIMIZE solve for an approximately optimal set of bids
3. Upload current bids to server

Prediction

Price Prediction

Competitive Equilibrium Prices

- prices at which supply = demand
  - all producers are profit-maximizing
  - all consumers are utility-maximizing

Notation

$\vec{p}$: price vector

$\vec{y}(\vec{p})$: cumulative supply of all producers

$\vec{x}(\vec{p})$: cumulative demand of all consumers

$\vec{z} = \vec{x} - \vec{y}$: excess demand

Price Prediction

Tâtonnement

$$\vec{p}_{t+1} = \vec{p}_t + \alpha_t \vec{z} \tag{1}$$

Walverine

$$\alpha_t \to 0 \tag{2}$$

SimAA

$$\vec{p}_{t+1} = \vec{p}_t + \alpha_t \max\{\vec{z}, 0\} \tag{3}$$

Price Prediction
2002 Finals

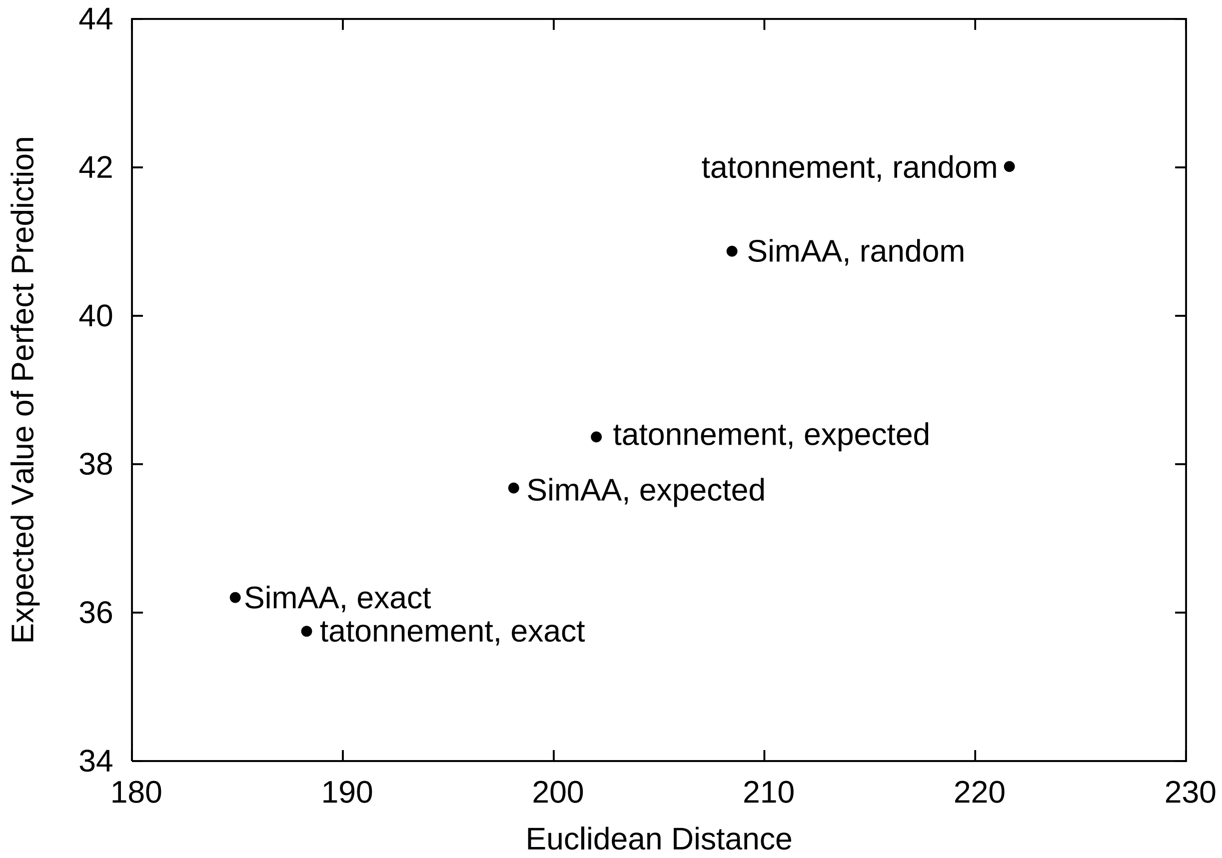

Price Prediction

2002 Finals [Wellman, et al. 04]

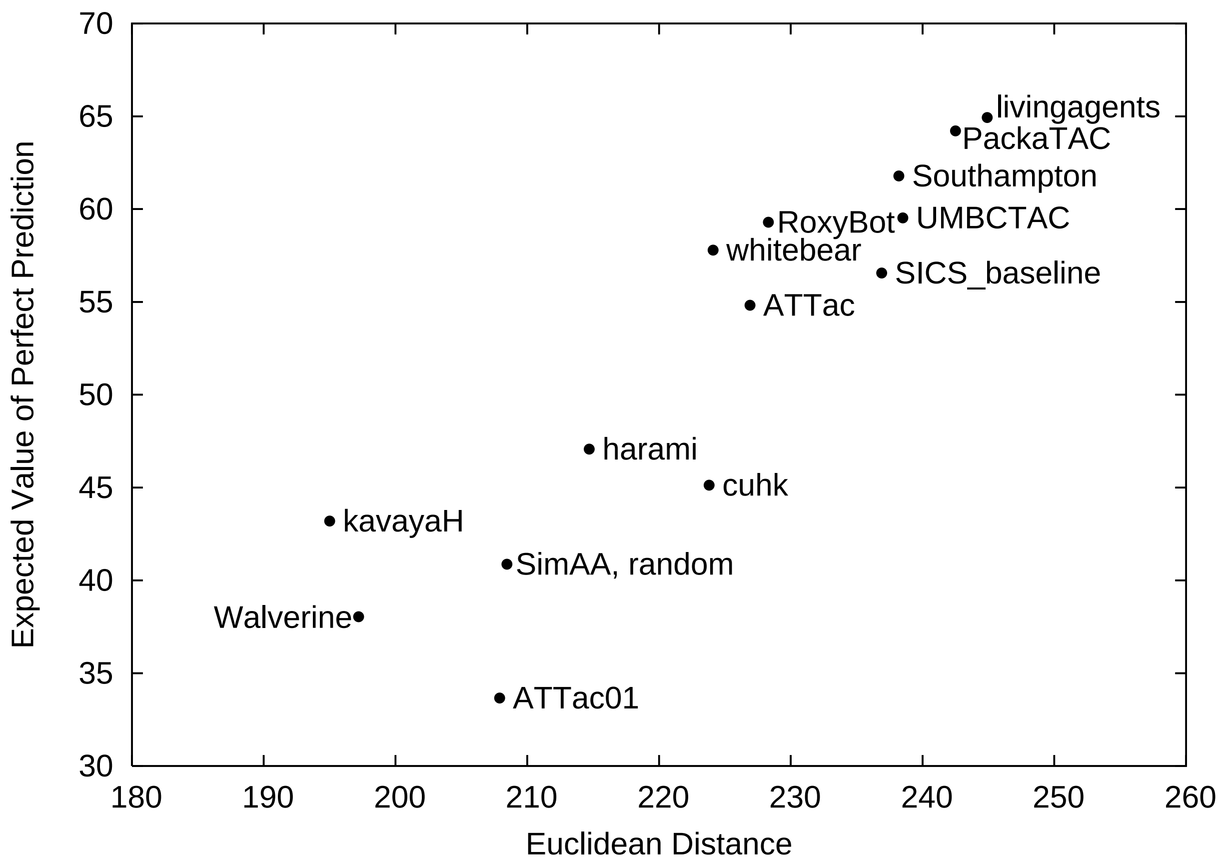

Bidding

Bidding in Pseudo-Auctions

Sealed-Bid & Second-Price

Pseudo-Auctions

- there is only one bidding agent
- clearing prices are specified by an exogenous model
- RULES
  - winner determination rule: win by bidding at least the clearing price
  - payment rule: pay the clearing price (second-price)

Bidding as Stochastic Optimization

What is the bidding agent's utility-maximizing set of bids?

Bidding in Sequential Pseudo-Auctions

Sealed-Bid & Second-Price

Given

an ordered set of goods $X = \{x_1, \ldots, x_n\}$

a combinatorial valuation function $v : 2^X \to \mathbb{R}$

a conditional CDF $F_j(b) = \Pr[p_j \leq b \mid p_1 \ldots, p_{j-1}]$

Bidding as an MDP

What is optimal bidding policy: that is,

what is the optimal bid to place on each good?

Markov Decision Process

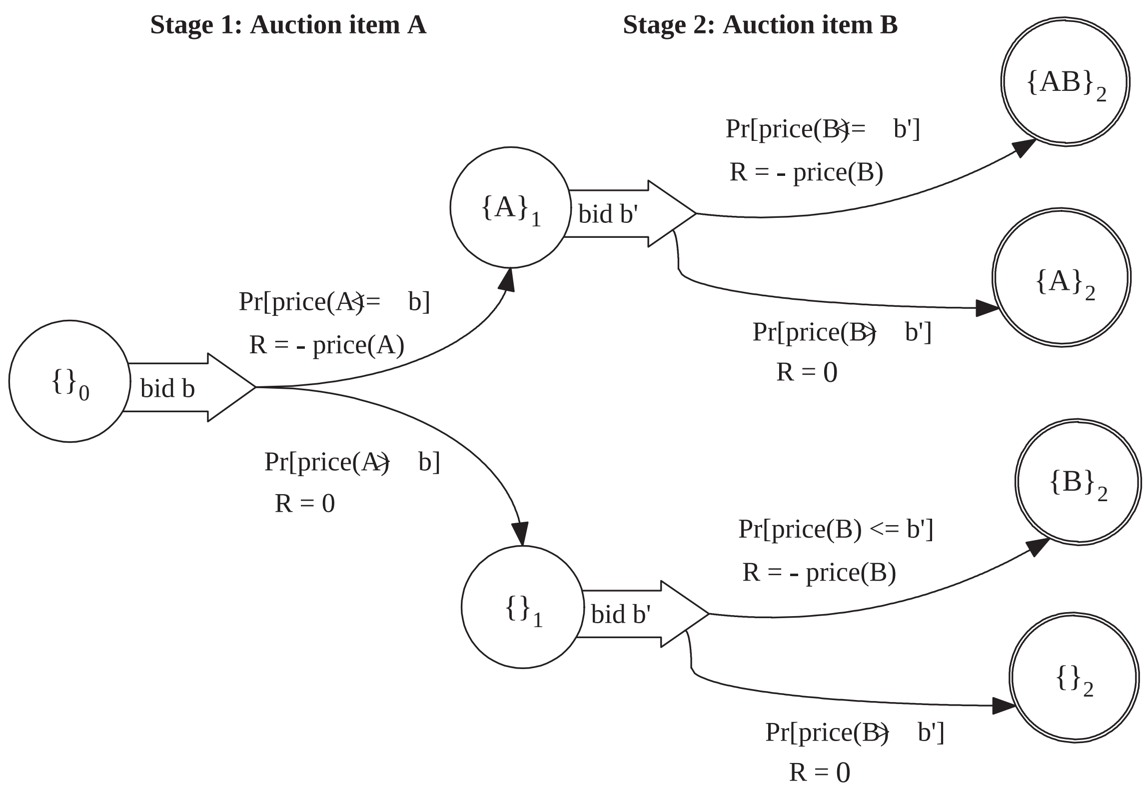

Bidding Marginal Utilities

Intuition

The marginal utility of a good x relative to a set of goods X is the difference between the utility of $X \cup \{x\}$ and the utility of $X \setminus \{x\}$.

Theorem [G and Boyan 2004]

Bidding marginal utilities is the unique optimal policy in the MDP that characterizes sealed-bid second-price sequential pseudo-auctions.

Bidding Marginal Utilities

Theorem

For all $1 \leq j \leq n$, at state $(Y, j-1)$,

$b_j^* = V((X \cup \{x_j\}), j) - V((X, j))$ is the optimal bid.

Proof

CASE $b_j < b_j^*$: If $b_j < p_j \leq b_j^*$, then you lose, but wish you hadn't, because:
$p_j \leq b_j^*$ implies $V((Y, j)) \leq V((Y \cup \{x_j\}, j)) - p_j$.

CASE $b_j > b_j^*$: If $b_j \geq p_j > b_j^*$, then you win, but wish you hadn't, because:
$b_j^* < p_j$ implies $V((Y \cup \{x_j\}, j)) - p_j < V((Y, j))$.

Bidding in Simultaneous Pseudo-Auctions
Sealed-Bid & Second-Price

Given

a finite set of goods X

a valuation function $v : 2^X \to \mathbb{R}$

an *additive* pricing function $p : X \to \mathbb{R}$

a distribution f over additive pricing functions

utility $u(Y) = v(Y) - p(Y)$, for all $Y \subseteq X$

Winner Determination Rule

$$x \in \text{Winnings}(X, p, b) \text{ if and only if } b(x) \geq p(x) \tag{4}$$

Bidding Problem

$$\text{SIM}(X, v, f) = \max_{b \in \mathbb{R}^X} \mathbb{E}_{p \sim f} \left[u(\text{Winnings}(X, p, b)) \right] \tag{5}$$

"Collapsing" Heuristics

Deterministic Bidding Problem

$$\text{DET}(X, v, p) = \max_{b \in \mathbb{R}^X} u(\text{Winnings}(X, p, b)) \tag{6}$$

Expected Value Method

$$\text{EVM}(X, v, f) = \text{DET}(X, v, \bar{p})$$

Implementation/Generalization

1. sample S scenarios $p_1, \ldots, p_S \sim f$
2. compute average prices $\bar{p} = \sum_{i=1}^{S} p_i$
3. run any "deterministic" bidding heuristic
   (i.e., assume sample mean prices are certain)

"Exploiting" Heuristics

(Stochastic) Bidding Problem

$$\mathrm{SIM}(X, v, f) = \max_{b \in \mathbb{R}^X} \mathbb{E}_{p \sim f} \left[u(\mathrm{Winnings}(X, p, b)) \right] \tag{7}$$

Sample Average Approximation

- sample S scenarios $p_1, \ldots, p_S \sim f$
- $\mathrm{SAA}(X, v, p_1, \ldots, p_S) = \max_{b \in \mathbb{R}^X} \sum_{i=1}^{S} u(\mathrm{Winnings}(X, p_i, b))$ (8)

Theorem [e.g., Ahmed and Shapiro 2002]
The probability that an optimal solution to $\mathrm{SAA}(X, v, p_1, \ldots, p_S)$ is an optimal solution to $\mathrm{SIM}(X, v, f)$ converges to 1 exponentially fast as $S \to \infty$.

"Collapsing" Heuristics

Example

$v(\text{camera} + \text{flash}) = 750$
$v(\text{camera}) = v(\text{flash}) = 0$

$p(\text{camera}) = 500$, with probability $\frac{1}{2}$
$p(\text{camera}) = 1000$, with probability $\frac{1}{2}$
$p(\text{flash}) = 50$, with probability 1

Predictions

$p(\text{camera}) = 750$, with probability 1
$p(\text{flash}) = 50$, with probability 1

Optimal Bid Vector

Bid Vector A: $(0, 0)$
Value$(A) = 0$

"Exploiting" Heuristics

Example

$v(\text{camera} + \text{flash}) = 750$
$v(\text{camera}) = v(\text{flash}) = 0$

$p(\text{camera}) = 500$, with probability $\frac{1}{2}$
$p(\text{camera}) = 1000$, with probability $\frac{1}{2}$
$p(\text{flash}) = 50$, with probability 1

Bid Vectors

Bid Vector A: $(0, 0)$ is optimal, with probability $\frac{1}{2}$
Bid Vector B: $(500, 50)$ is optimal, with probability $\frac{1}{2}$

Value of Stochastic Information = 75

$\text{Value}(B) = \frac{1}{2}(200) + \frac{1}{2}(-50) = 75$
$\text{Value}(A) = 0$

“Collaping” Heuristic:

Straight Marginal Utility

$$\left(\max_{Y\subseteq X\setminus\{x\}}\left[v(Y\cup\{x\})-\frac{1}{S}\sum_{i=1}^{S}p_i(Y)\right]\right)-\left(\max_{Y\subseteq X\setminus\{x\}}\left[v(Y)-\frac{1}{S}\sum_{i=1}^{S}p_i(Y)\right]\right) \quad (9)$$

“Collaping” Heuristic:

Straight Marginal Utility

$$\left(\max_{Y\subseteq X\setminus\{x\}}\left[v(Y\cup\{x\}) - \frac{1}{S}\sum_{i=1}^{S} p_i(Y)\right]\right) - \left(\max_{Y\subseteq X\setminus\{x\}}\left[v(Y) - \frac{1}{S}\sum_{i=1}^{S} p_i(Y)\right]\right) \quad (9)$$

“Exploiting” Heuristic:

Average Marginal Utility

$$\frac{1}{S}\sum_{i=1}^{S}\left(\max_{Y\subseteq X\setminus\{x\}}[v(Y\cup\{x\}) - p_i(Y)] - \max_{Y\subseteq X\setminus\{x\}}[v(Y) - p_i(Y)]\right) \quad (10)$$

[Stone, et al. 01]

Test Suite of Bidding Heuristics

“Collapsing” Heuristics

- RoxyBot–2000 and RoxyBot–2000*
- Straight Marginal Utility

“Exploiting” Heuristics

- RoxyBot–2002 and RoxyBot–2002*
- Average Marginal Utility
- SAA and SAA* (RoxyBot–2006)

IV. Experiments

Offline Experiments

- TAC Travel Pseudo-Auctions
  - Simultaneous: Flights and Hotels
  - Multiunit Sequential: Flights and Hotels

Online Experiments

- Online: TAC Travel Games (“Real” Auctions)

TAC Travel Online Experiments

Scores, Utilities, and Costs

Rank	Agent	Score	Utility	Cost
1	SAA*	2707	8268	5561
2	SAA	2678	8257	5579
3	2002	2632	8250	5618
4	2002*	2627	8401	5774
5	2000*	2622	8399	5777
6	2000	2620	8307	5687
7	AMU	2521	8308	5787
8	SMU	2511	8336	5825

TAC Travel Online Experiments

Scores, Utilities, and Costs

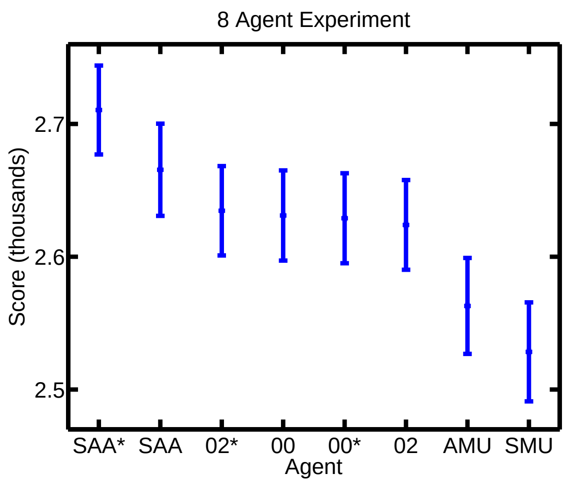

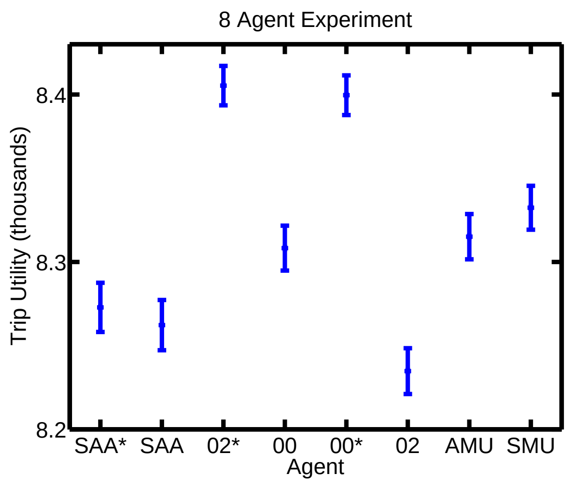

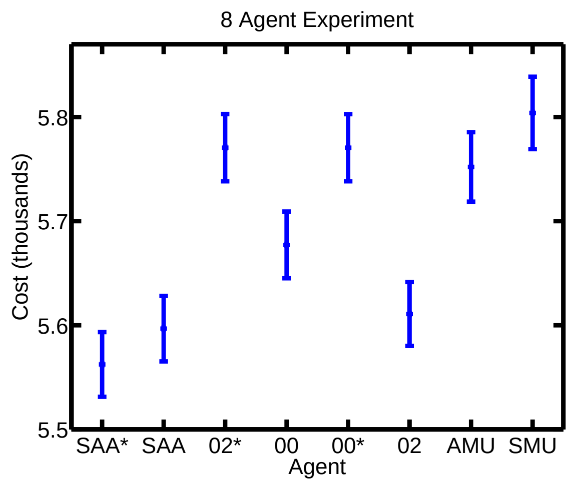

TAC Travel Online Experiments

Parameter Settings and Runtimes*

Agent	E	S	P	# of Optimizations	SG	BC	Total
2000	–	50	–	$2n+1$	9.4	1.0	10.4
2000*	–	50	–	$2n+1$	9.0	1.1	10.1
2002	15	–	25	$(2n+1+P)E$	7.0	5.3	12.3
2002*	15	–	25	$(2n+1+P)E$	7.0	4.7	11.7
AMU	–	15	–	$2nS$	2.3	10.2	12.5
SMU	–	50	–	$2n$	8.7	1.5	10.2
SAA	–	50	–	N/A	8.8	1.7	10.5
SAA*	–	50	–	N/A	9.0	1.6	10.6

*Machines were not dedicated.

Summary

Theory

- Sequential Pseudo-Auctions
  1. Bidding Marginal Utilities is an Optimal Policy
- Simultaneous Pseudo-Auctions
  1. Bidding Marginal Utilities is NOT an Optimal Policy
  2. Sample Average Approximation Approximates an Optimal Policy

Experiments

- SAA performs well in simultaneous auctions
- MU performs well in multiunit sequential auctions
- SAA* performed well (WON!) TAC Travel 2006 ("real" auctions)

Future Directions

Optimal	Simultaneous	Sequential
Deterministic	RoxyBot–2000	EASY
Stochastic	SAA, as $S \to \infty$	Marginal Utility

Practical	Simultaneous	Sequential
"Collapsing"	EVM, RoxyBot–2000	Straight MU
"Exploiting"	SAA, RoxyBot–2002	Average MU

Given this set of bidders, what is the preferred auction design?

- from the point of view of the auctioneer
- from the point of view of the bidders

Thank You!

Amy Greenwald
amy@brown.edu
http://www.cs.brown.edu/people/amy

http://www.sics.se/tac