Adaptive Market Design - The SHMart Approach

Harivardan Jayaraman
hari5851@cs.utexas.edu

Sainath Shenoy
sainath@cs.utexas.edu

Department of Computer Sciences
The University of Texas at Austin

Abstract

Markets are essentially important in the trading of any products. Traditionally markets have been regulated by humans who maintain the fees levied on its traders. With the advent of autonomous agent technologies, new strategies to automate this process can be discovered. This paper introduces SHMart, an agent deployed in the testbed of CAT (an autonomous market scenario containing autonomous traders), and presents its intelligent strategies through analysis of controlled experiments.

Introduction

Markets are omnipresent during the trading of trading of any complex product. Traditionally markets have been created by humans who maintain the market's fees and the profits made by the traders. With the foray of autonomous agents in our day to day life, there has been recent interest in the process of automating these market mechanisms.

The process of creating an autonomous agent is quite complex since various different priorities need to be maintained by an agent at any given time. It should try to maximize its profit, keeping in mind simultaneously, the profits made by the other competing markets as well as social welfare of the traders.

The CAT Mechanism Design scenario provides such a competitive environment where independently developed agents can be deployed and tested against each other over the course of many simulations of an actual market.

The CAT Market Mechanism Scenario

In this section we describe a basic overview of the CAT domain. In a CAT game, the agents act as *specialists* in a simulated market managed by a game server. The autonomous *traders* are divided into *buyers* and *sellers* who exchange goods in such a simulated market environment. Each CAT game consists of a number of *days* and each day consists of a number of *rounds*. The specialist with the maximum profits wins.

Market Architecture

The market consists of specialists and traders. At the start of every trading day, the traders' private values are initialized and sellers are each endowed with one unit of a homogeneous good. The traders buy or sell based on whether they are a buyer or a seller by placing bids or asks at a particular specialist. The goal of every trader is to maximize a utility function based on the demand and supply in the market. The specialist controls the market by matching buyers with sellers and charges traders for the activity in its market.

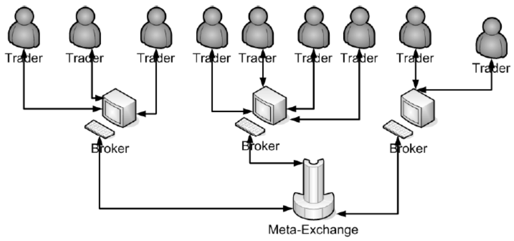

Fig.1. CAT Market Architecture

Each trader has a demand and supply function, a bidding strategy and a specialist selection strategy.

Bidding Strategies

The traders are essentially intelligent and adaptive and learn from past transactions. Two types of trader strategies (ZIP [1], GD [4]) were chosen for experimental purposes.

Specialist Selection Strategy

The traders have an explorative phase and an exploitation phase. The traders compute a confidence value for each specialist during the explorative phase where specialists are chosen using a multi-armed bandit [9] or a StimuliResponse [10] approach.

In this paper we describe SHMart, an adaptive agent deployed in the CAT testbed. We describe its various intelligent components and strategies and how they were combined together to create an effective market mechanism in the midst of fierce competition.

We start off with the description of the CAT strategy space and possible options available to a Specialist designer for the CAT Tournament. We then describe the price formation mechanism and the implementation of the Preston MacAfee Double Auction [3] Mechanism implementation in SHMart. We also look at the fee calculation mechanism in that uses agent reward estimation. The last 2 sections showcase experimental results and future work on improving the current design.

CAT Strategy Space

A Specialist designer has a somewhat limited strategy space in CAT. There are essentially four areas for innovation viz.:

1) *The Charging Policy:* This policy is responsible for the fees charged by a specialist. The starter strategy available here is a Fixed Charging Policy. Intuitively this seems to be the prime factor for specialist design. Each specialist can make profits by charging each trader for:

1. *Registration:* A one time fee paid by an agent at the start of the trading day.
2. *Transactions:* A fee paid whenever an agent transacts at a given market.
3. *Shouts placed:* A fee paid whenever an agent places a shout (a bid or ask).
4. *Information requested:* A subscription fee paid by and agent to get market order book information.
5. *Profit margin fee:* The percentage cut of profit made by an agent during a transaction.

Most trader agent implementations maintain a reward vector for each specialist that reacts to profits made by an agent while transacting with it. As a result high prices may increase specialist profit, but only temporarily and usually results in loss of traders in subsequent days. Similarly, low prices attract high volumes of traders, but may require several days to make substantial profits.

2) *The Clearing Policy:* This policy determines when and how to clear an auction. The when part of the policy checks a certain condition to start the process of clearing, while the how part defines it. The starter strategies available here are Round Clearing and Probabilistic Clearing. Both the strate-

gies use a simple double auction clearing strategy that sorts the bids in ascending order and asks in descending order and matches them in order. The clearing policy plays a major role in price formation in the mechanism.

3) *The Pricing Policy:* This policy determines the price at which to clear a transaction given the bid and ask. A discriminatory pricing condition is provided as starter strategy that biases the price toward either the bid or the ask, based on the value of k, the bias [7]. The pricing policy is another important component in price formation and can be used effectively to improve bidder/seller starved conditions.

4) *The Shout Acceptance Policy:* This policy represents the shout improvement rule for the double auction, enabling faster price formation [8]. The NYSE, market improvement rule is provided as a starter strategy.

SHMart Strategies

We suspect that most specialist designers would concentrate on the charging policy design. For SHMart, however, we preferred to take a more balanced and robust approach and concentrated on all four policies. We divided the Specialist design task into two modules: Price Formation and Fee Calculation.

1) Price formation Strategy

Clearing Policy:

The clearing policy that we implement is a hybrid between the round clearing and the probabilistic clearing policies: we clear the auction, matching any outstanding bids and asks on a round basis using the Preston MacAfee Double Auction [2][3][8], while in the last round we use a generic double auction [6] that clears probabilistically every time a new shout is placed.

The Preston MacAfee Double auction is a truth revealing dominant strategy mechanism [3] for both buyers and sellers. The mechanism can be illustrated as follows:

1) Rank the bids in descending order $b_1>b_2>b_3\dots$.
2) Rank the asks in ascending order $a_1<a_2<a_3\dots\dots$
3) Find the efficient trade quantity k such that $b_k>=a_k$ and $b_{k+1}<a_{k+1}$.
4) Calculate price p_0 for all k trades as $0.5*(b_{k+1}+a_{k+1})$
5) If($p_0>b_k$ or $p_0<a_k$) trade the first k-1 bids at b_k and the first k-1 asks at a_k (losing out on b_k and a_k).
6) Else trade the first k bids and asks at p_0.

Due to the absence of a dual priced transaction scheme, we implemented step 5 to trade all k bids and asks at $0.5*(a_k + b_k)$. The mechanism, we assume, still retains its truth revealing dominant strategy, without any loss of trades.

The absence of strategic behavior means that the properties of the equilibrium can be established purely on the characteristics of the underlying distribution, without reference to bidding behavior. At most one transaction is lost (which is least valuable), thus the per trader efficiency loss is of the order of 1/n [3][5]. Lastly both the mechanism and equilibrium strategies can be defined without reference to the underlying distributions, so that game and equilibrium allocations are not sensitive to bidder types. Thus this results in an essentially bid/ask independent mechanism [2].

Shout acceptance:

The NYSE, market improvement rule enables faster price formation, but results in lower shout fee collections. We hypothesize that using the PMDA described above, would speed up price formation and enabling the mechanism to accept all shouts will drive up shout fee collections.

Discriminatory Pricing

The starter strategy maintains the k price bias at 0.5 [7], allowing maximum efficiency. We hypothesize that in a market that is either bidder/seller starved or has an imbalance between bids and asks, symmetrical prices will not better the situation, or worsen it further. We use a minority biased discriminatory pricing scheme that biases the price toward the ask if bids are in minority and toward the bid if asks are in minority. We expect that this biasing, would encourage order book and clearing price reactive agents to rectify the bid/ask starved situation.

$$\text{Clearing price} = k*\text{Clearing bid} + (1-k)*\text{Clearing ask} \quad (1)$$

where

$$k = \frac{\text{Pending bids}}{(\text{Pending bids}+\text{Pending asks})} \quad (2)$$

If there are less pending asks as compared to bids (seller starved market), we must price closer to the bid, so as to give sellers greater profit and enable them to increase their numbers and vice versa.

2) Fee Calculations:

The market selection behavior of traders can be classified into explorative and exploitative phases. Typically, in the initial days of any game, traders explore and test every market available and form a reward vector for each market. Based on these reward vectors they choose a market and conduct trades in it. Owing to this feature of traders, we divided our specialist strategy to work differently in the different phases.

In the explorative phase, SHMart charges lower than every other specialist in the market for each of the fees. This ensures that every trader forms a high reward vector for it. Simultaneously, SHMart keeps note of the number of traders that have entered each market and computes the average of this value to know when to come out of the explorative phase.

Every trader has a *specialistReward* map which maintains the estimated reward a specialist gives to the trader. The trader reward calculation also calculates the specialist profit as a by product. The accuracy of this estimated profit is quite high, and is quite close to the actual profit for each specialist.

We hypothesize that the incremental revenue earned by a specialist due to a trader is as follows:

Incremental specialist reward =

Registration fee +

Information Fee {if information seeking} +

(Avg. Transaction Probability*Transaction Fee)+

((Avg. bid ask spread/2)*
(Avg. transaction probability*Profit fee))+

(Average shouts per agent per day*shout fees) (3)

Assuming that agents bid their true value,

Agent Day Reward =

((Avg. bid ask spread/2)*
Avg. transaction probability) –

(Incremental specialist reward) (4)

All averages are calculated using moving averages based on data seen in our market. We hypothesize that these averages hold in all markets.

This reward mechanism can also be used to predict the traders that would be present in our market (or in any other market) on the next day, based on which market has a higher reward value for a trader. This mechanism is self correcting and looks at next day actual results and adjusts reward values accordingly. Also between game iterations we look at actual revenues made by specialists to weigh equation (3) for each specialist.

SHMart focuses only those specialists that are above the average estimated profit and tries to lure the traders of the specialist with the maximum traders by pricing lower than it.

Hence it gets the best of the profit as well as the number of traders. Since SHMart already has a good reward level for each trader, its trader count increases.

Once it becomes the specialist with the highest profit, it strives to stay that way, and hence prices low, keeping its traders satisfied. It charges higher only if it detects that the specialist with the second highest trader count is charging more than itself. In such a case, it increases its price by an amount which still ensures that it remains below the other specialist. The change in price is dictated by the trader count history of two days and the predicted trader count of one day.

Experimental results

We performed a number of experiments while tuning our strategy. We also developed several dummy specialists to model possible competition. In this section we look at two important results:

1) Decrease in shout shading due to use of PMDA vs. GDA.

The anytime clearing scheme of the GDA presented to us a possible improvement area. We believe that this scheme gives no incentive for traders to bid their true value, and results in under reporting or shout shading. The PMDA on the other hand is a truth revealing mechanism and we expected it to reduce shout shading. We ran tests with 2 markets one using a GDA and the other a PMDA. A trader mix of 109 sellers and 52 buyers was used with ZIP learning strategies. Fig.2 shows the average shout shading per day.

In most of the tests PMDA resulted in lesser or equal shading when compared to GDA but in no game was GDA found to induce traders to shade less. The PMDA results in better value reporting shouts as compared to a GDA.

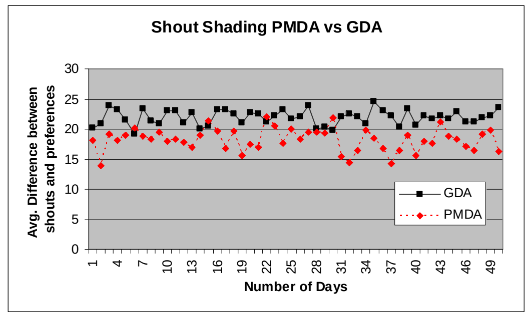

Fig.2. Shout Shading PMDA vs. GDA

2) SHMart vs. SHMartRamp

SHMartRamp is a specialist that we created from SHMart which differs in its exploitative phase fee calculation mechanism. It uses only a trader prediction mechanism and ramps up its prices when it predicts it will have more traders.

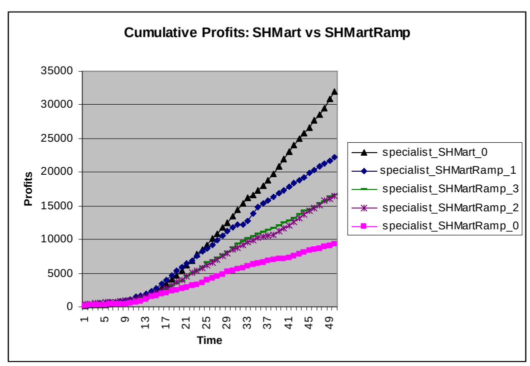

Fig.3 SHMart vs. SHMartRamp

We performed tests against 4 SHMartRamp agents for a trader mix of GD and ZIP. The cumulative profit curve for one such game is as shown above (Fig. 3).

2) SHMart vs. Fixed

We tested our specialist against the Fixed charger dummy specialist. In our results we used a combination of ZIP and GD strategy traders.

The Fixed charging specialists charged a constant price throughout the game. We deployed SHMart against 4 other specialists, all charging fixed prices with high disparity, for 3 games each of 50 days with 10 rounds. The aggregate result of these games is as shown in Fig.4.

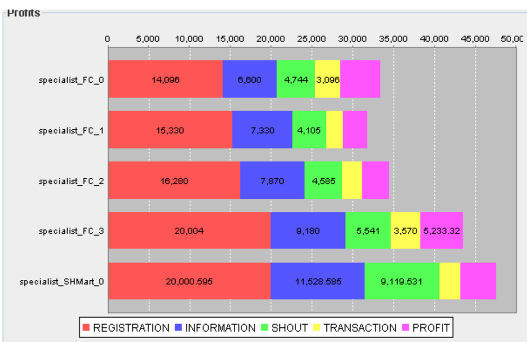

Fig.4 SHMart vs. Fixed Aggregate profits

3) SHMart vs. Random

The Random charging specialists initially selects a price and keep perturbing randomly with a 20% variance. This ensures that atleast one specialist in each game finds the optimal charges. We played SHMart with 4 Random chargers for 10 games each consisting of 50 days with 10 rounds. The aggregate results for these games are as shown in Fig. 5 and the cumulative profit of one such game has been shown in Fig. 6.

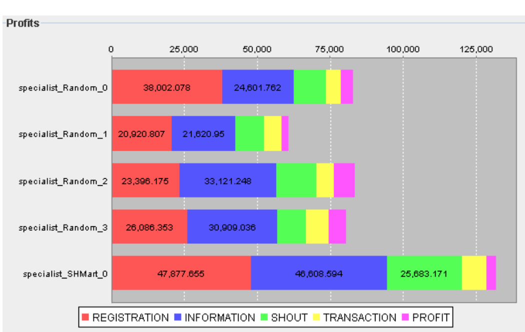

Fig. 5 SHMart vs. Random Aggregate profits

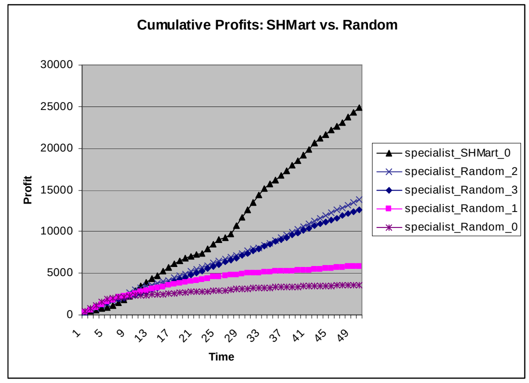

Fig. 6 SHMart vs. Random

We also tested SHMart against a combination of these three types of specialists for similar games. We used one SHMartRamp, one Fixed charging and two Random charging specialists. The results are as shown.

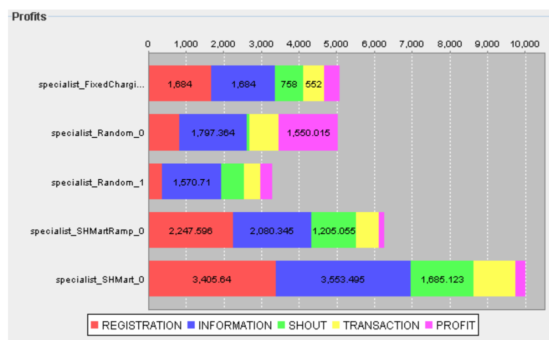

Fig. 7 SHMart vs. Various Strategies

Future work

The CAT testbed for automated mechanism design is a relatively new research area in autonomous agent research. A lot of the strategies we developed were therefore empirical or specific in nature. We identify several improvement areas in our specialist through SHMartPlus a hybrid of the strategies of SHMart and SHMartRamp.

The SHMartRamp specialist attempts to predict the next day trader behavior based on current day prices. As part of the initial implementation of SHMartPLus, we had developed a next day profit estimator strategy, but could not tune it in time for the competition. We believe that this mechanism can help determine prices better and make SHMart more robust.

SHMartPlus also looks at past game data to determine initial prices as well as gather better profit estimates.

The SHMartRamp strategy circumvents alternate low-high pricing strategies or scenarios where there is a high price disparity that benefits high pricing specialists. This strategy can be improved to adapt to other scenarios as well.

Conclusion

The paper discusses SHMart, an adaptive market mechanism. In addition, it also analyzes the performance of SHMart against different types of specialists in the CAT testbed. The fine tuning and improvement of the adaptive and predictive components remains important areas for future work.

References

1) Cliff, Dave, and Janet Bruten, "Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets", Tech. Rep. no. HPL-97-141, Hewlett-Packard Laboratories, (1997).

2) K. Deshmukh, A. Goldberg, J. Hartline, and A. Karlin, "Truthful and Competitive Double Auctions." In Proc. 10th European Symposium on Algorithms. Springer--Verlag, 2002.

3) McAfee, R. Preston, 1992. "A dominant strategy double auction," Journal of Economic Theory, Elsevier, vol. 56(2), pages 434-450, April.

4) S. Gjerstad, and J. Dickhaut, "Price Formation in Double Auctions", Games and Economic Behavior, 22, pp. 1-29 1998.

5) D. K. Gode, and S. Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality", Journal of Political Economy, Feb 1993, 101,1, pp. 119-137.

6) D. Friedman. "The double auction institution: A survey." In D. Friedman and J. Rust, editors, Santa Fe Institute Studies in the Sciences of Complexity,chapter 1, pages 3{25. Perseus Publishing, Cambridge, MA, 1993.

7) M. Satterthwaite and S. R. Williams. "The Bayesian theory of the k-double auction." In D. Friedman and J. Rust, editors, The Double Auction Market: Institutions, Theories and Evidence, Santa Fe Institute Studies in the Sciences of Complexity, chapter 4, pages 99{123. Perseus Publishing, Cambridge, MA, 1993.

8) Jonathan Bredin and David C. Parkes. "Models for truthful online double auctions." In Proc. 21st Conference on Uncertainty in Artificial Intelligence (UAI'2005), pages 50–59, 2005.

9) P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "Gambling in a rigged casino: the adversarial multi-armed bandit problem," in Proc. 36th Annual Symposium on Foundations of Computer Science, Los Alamitos, CA: IEEE Computer Society Press, 1995, pp. 322--331.

10) Gibson, F. P. 2002. "Is It Better to Forget? Stimulus-Response, Prediction, and the Weight of Past Experience in a Fast-Paced Bargaining Task." Comput. Math. Organ. Theory 8, 1 (May. 2002), 31-47.