Evolving Market Design in Zero-Intelligence Trader Markets

Vibhu Walia, Andrew Byde, Dave Cliff
Intelligent Enterprise Technologies Laboratory
HP Laboratories Bristol
HPL-2002-290
October 30th , 2002*

Email: walia@hp.com, dave_cliff@hp.com, andrew_byde@hp.com

CDA, automated mechanism design, ZI-C Agents, hybrid markets, agent-based simulation, evolutionary strategy, BICAS

The Continuous Double Auction (CDA) is one of the most popular of all auction markets in the world. Some of the biggest trade markets in the world including the New York Stock Exchange (NYSE) and the Chicago Mercantile Exchange are organised as CDAs. Previous work by Cliff has shown that previously unknown 'hybrid' variants of the traditional CDA can be found to give the most desirable market dynamics. Cliff's results were based on experiments conducted using a computational simulation of the CDA populated by electronic Zero Intelligence Plus (ZIP) traders and his work uses a GA to co-evolve the market mechanism as well as the ZIP agent parameters. The exclusive use of the ZIP trading algorithm in all his work raises questions about the robustness of his results to any change in the trading algorithm used. In this thesis we use a self-adaptive Evolutionary Strategy (ES) to explore the space of possible auction types in a CDA populated by Gode and Sunder's cognitively simple Zero Intelligence (ZI) traders. We show that hybrid CDAs are still preferred over traditional variants and our results provide the first demonstration that hybrid variants of the CDA can provide favourable dynamics for trading strategies other than ZIP.

* Internal Accession Date Only
Approved for External Publication

Evolving Market Design in Zero-Intelligence Trader Markets

Vibhu Walia
Hewlett-Packard Laboratories
Filton Road, Bristol BS34 8QZ, England, U.K.
+44 (0) 117 312 9110
walia@hp.com

Andrew Byde[1]
Hewlett-Packard Laboratories
Filton Road, Bristol BS34 8QZ, England, U.K.
+44 (0) 117 312 8764
dave_cliff@hp.com

Dave Cliff
Hewlett-Packard Laboratories
Filton Road, Bristol BS34 8QZ, England, U.K.
+44 (0) 117 312 8189
andrew_byde@hp.com

ABSTRACT

The Continuous Double Auction (CDA) is one of the most popular of all auction market-mechanisms in the world. Some of the biggest trade markets in the world including the New York Stock Exchange (NYSE) and the Chicago Mercantile Exchange are organised as CDAs. Previous work using Genetic Algorithms (GAs) for automated mechanism design by Cliff has shown that previously unknown 'hybrid' variants of the traditional CDA can lead to the most desirable market dynamics. Cliff's results were based on experiments conducted using a computational simulation of the CDA populated by electronic Zero Intelligence Plus (ZIP) traders and his work uses a GA to co-evolve the market mechanism as well as the ZIP agent parameters. The exclusive use of the ZIP trading algorithm in all his work raises questions about the robustness of his results to any change in the trading algorithm used. In this paper we use a self-adaptive Evolutionary Strategy (ES) to explore the space of possible auction types in a CDA populated by Gode and Sunder's cognitively simple Zero Intelligence Constrained (ZI-C) traders. We show that hybrid CDAs are still preferred over traditional variants and our results provide the first demonstration that hybrid variants of the CDA can provide favourable dynamics for trading strategies other than ZIP.

Keywords

CDA, automated mechanism design, ZI-C agents, hybrid markets, agent-based simulation, evolutionary strategy.

1. INTRODUCTION

Although the Continuous Double Auction (CDA) is one of the most popular types of market-mechanisms in the world today, it is also one of the least understood. Its complex nature means that game-theoretic analysis of the CDA is intractable except for the simplest of cases. Experimental methods have long been used to explore the mechanisms behind some of its most curious properties and CDA experiments with human participants have been organised since the early 60's [12].

Agent-based simulations offer us a powerful tool that we can use to study the properties of a CDA using empirical methods. Agent-based methods are not only convenient and economical but also offer a tightly controlled environment in which a diverse range of experiments can be performed.

The Double Auction is actually a general name for a broad class of trading institutions in which both buyers and sellers can submit bids and offers. This is as opposed to only buyers shouting offers (as in an English Auction), or only sellers shouting bids (as in a Dutch Auction). Most stock and commodity markets around the world including those at New York, Tokyo, Frankfurt and Chicago are organized as variants of the basic CDA mechanism.

Some recent work on automated market mechanism design using GAs by Cliff [3,4] have shown that hybrid variants of the CDA, which are unlike previously known versions, can give the most desirable market dynamics. Work by Phelps et al [9] uses a Genetic Programming (GP) based approach to co-evolve a trade-settling formula for the CDA along with the trading algorithm parameters. These attempts at automatic market design are particularly relevant in the context of market mechanism design for electronic trading agents. Electronic agents are willing participants in any kind of auction and there is nothing that stops us from adopting these 'peculiar' hybrid versions of the CDA for online auction markets. However, for wide adoption of these markets we need to conclusively prove that: -

1. Automatically designed markets consistently provide larger gains compared to more traditional market mechanism types.
2. These gains are achieved regardless of the trading strategy used.

Previous work has addressed the first of these concerns but has failed to address the second.

[1]Address all correspondence to this author.

In this work we have used a self-adaptive Evolutionary Strategy to search through the space of possible auction types, where the market-mechanism is parameterised by a real-valued parameter Q_s. This parameter was first introduced by Cliff [3] and forms the basis for his work exploring automatic market design [3,4]. We have used the Zero Intelligence Constrained (ZI-C) traders, first introduced by Gode and Sunder [9], to populate our simulations of the CDA. ZI-C traders submit random bids subject to the constraint that they cannot shout loss-making offers and specify the lower bound on the cognitive ability that is needed for a trader participating in a CDA. We also incorporate some suggestions made by Preist and van Tol [11] in the methodology we use to organise the CDA.

Our results demonstrate that 'hybrid' variants of the CDA dominate more traditional variants in many cases. We go on to explore the possible reasons for the superiority of the new 'hybrid' variants over more conventional forms of the CDA. We conduct experiments to optimise Q_s values over a two-dimensional space of test-markets parameterised by the slopes of the supply and demand curves. We plot three-dimensional landscapes of optimum Q_s values for all markets in our test space and discuss possible reasons for the regularities observed.

The rest of the paper is organised as follows. Section 2 discusses the previous work on Q_s based CDAs and puts our work in context. Section 3 describes our experiments, in which we use an ES to find the optimum value of Q_s for any market organised as a CDA. In Section 4 we try to investigate what underlying characteristics of a given market could lead to a particular Q_s value being optimal. Finally, Section 5 provides a summary of our work and points out some directions in which this work can be extended.

2. BACKGROUND WORK

2.1 ZI-C Traders

Zero Intelligence Constrained or ZI-C traders were designed by Gode and Sunder (1993) to study the lower limit of rationality required to participate in a Double Auction market. They submit, "…random bids and offers distributed independently, identically and uniformly over the entire feasible range of prices" subject to a budget-constraint and thus cannot quote loss-making bids or offers. There are two reasons why we have chosen ZI-C agents for our experiments. Firstly, experimental results obtained with ZI-C trader experiments allow us to make a strong statement about result-validity and robustness. Favourable results obtained with ZI-C traders indicate that the market mechanism used is incentive-compatible and probably robust to changes in the trading strategy used. Secondly, the original motivation behind the design on the ZI-C agent strategy was to address the assumption, common in Economics literature but invalid in practice, that the behaviour of all traders (agents) is perfectly rational. Since ZI-Cs represent the lowest end of the spectrum in terms of trader-rationality, it follows that any results obtained in ZI-C markets are robust to trader-irrationality. This is a fundamental concern in the field of market-mechanism design and market-mechanisms that perform well despite trader irrationality are naturally preferred over those mechanisms that deliver high-performance only with perfectly rational traders.

The simplicity of the ZI-C trading algorithm, which makes it so appealing, also leads to a large variance in test results. For this reason in all our experiments we have averaged the fitness measurements over many ($\sim 10^3$) market trials. There has been a strong tradition of carrying out market trials in a staggered manner with trading taking place over continuous trading periods or days [2, 3, 4, 9, 12]. The chief motive for this was to allow time for learning and a resulting convergence towards equilibrium price. Since the behaviour of ZI-C traders is uniformly random it does not make sense to use trading days in ZI-C markets.

2.2 Persistent Shout Double Auction

In implementing the bidding process we have incorporated two suggestions made independently by Preist and van Tol [11] and Das et al [6]. Preist and van Tol studied a variant of the CDA called the CDA with an order-queue, which they termed a *persistent shout double auction market*. In a CDA with an order-queue a trader's current bid or offer persists until the trader makes another or is able to execute a trade at that price. Preist and van Tol [11] demonstrate that agents in a CDA with an order-queue reach equilibrium much faster, maintain a more stable equilibrium and are more robust to changes in learning rate (for ZIP agents). The New York Stock exchange is organized as a CDA with an order queue but has an additional constraint that any new bids or offers must improve on the existing ones. This constraint is commonly referred to as the *NYSE rule* in trading-agent literature. Thus, we have implemented a market that is a CDA with an order-queue but with the additional *NYSE rule* constraint [11, p. 6].

2.3 Automating Market Design

One way to automate the process of market design is to parameterise the description of a given market type. Once we have a complete parametrisation of a market mechanism, the design process can be seen as a parameter optimisation problem. This methodology was used by Cliff [3]. Cliff had previously used a GA to optimise the parameters for his ZIP trading algorithm in [2]. In [3] he added a single parameter Q_s to the tuple of ZIP initialising parameter values. In an ordinary CDA, any buyer or seller can shout a bid or an offer at any given time. The modification Cliff proposed to this mechanism has to do with choosing a balance between buyer and seller shouts. Q_s denotes the probability of the next shout coming from a seller. Thus, in an ordinary CDA Q_s=0.5. Even though it is straightforward to implement this modified mechanism in a computer simulation of a CDA hybrid values of Q_s (i.e. not equal to 1, 0 or 0.5) are unlike any known types of auction markets.

Cliff [3] reports that in three tests that he conducted on three different markets M1, M2 and M3, the optimum value of Q_s that is evolved using a GA search are 0.0001, 0.07 and 0.16. In a subsequent work [4], he describes the results on test market in which the supply and demand schedule is changed suddenly between experiments. This is referred to as a *market shock* and seeks to mirror the occurrence in everyday financial market when a sudden event can lead to a massive change in trader preferences, thereby changing the underlying supply and demand curves significantly. In these experiments M1M2 and M2M1, it was found that the optimum Q_s were 0.25 and 0.45 (statistically indistinguishable from 0.5).

However, there are several questions that could be raised about the validity of Cliff's results. Firstly, all experiments were performed with ZIP traders with the initialisation parameter values for the ZIP algorithm co-evolving alongside. Hence the exclusive use of the ZIP trading algorithm raises serious questions about the generalisation and robustness of these results to a change in the trading algorithm used. Secondly, Cliff uses Smith's 'coefficient of convergence' α to measure the performance of the market mechanisms being tested. We now know that high performance on one market-performance metric does not guarantee performance on another. Do Cliff's results generalise well to other fitness measures?

We seek to answer both these questions in our work. Firstly, we have used ZI-C traders in our work, which are not only a different trading-algorithm from ZIP but also define the lower bound on rationality required to participate in the CDA. Secondly, we have used allocative efficiency as a measure of market fitness. It has been shown that ZI-C traders can extract a large percentage of the available surplus (~97%) in a market [9]. Moreover, we believe that allocative efficiency is a more direct measure of the motives behind market design. Faster and more stable convergence do not mean much to a trader or to a market organiser on its own. It must be accompanied by other measures like an increase in average profit (from the point of view of a trader) or an average increase in the total profit made in the market or the percentage of viable trades executed (from the point of view of a market organiser). Allocative efficiency directly reports on the latter (and profit dispersion on the former).

3. EVOLUTIONARY DESIGN OPTIMISATION

3.1 Algorithm

We have used a self-adaptive Evolutionary Strategy (ES) style approach to design an Evolutionary Algorithm to optimise the value of Q_s for a given market. Evolutionary Strategies are the application of the Evolutionary Computation (EC) paradigm to real valued parameter optimisation [13,14]. ES were designed to be self-adaptive from inception and studies by Fogel [7,8], and Back and Schwefel [1] have shown that an ES with a self-adaptive mutation operator usually perform better than an ES without a self-adaptive operator on the same problems. A general self-adaptive strategy, as outlined by Yao, Liu and Lin [17], is shown in Figure 1. Figure 2 shows pseudo-code for the evolutionary algorithm that we have designed to optimise Q_s values for a given market.

3.1.1 *Test Cases*

The test cases MZIC-I and MZIC-II have 6 buyers and sellers each. The supply and demand curves for each of these markets are shown below in Figures 3 and 4.

halting criterion satisfied else go back to

Figure 1. Schematic description of a typical Self-Adaptive Evolutionary Strategy (ES) or Evolutionary Programming (EP) style approach

1) Initialise population $\mathbf{Q_s[\mu]}$ with *U[0,1]* where *U[0,1]* is a uniformly distributed random number between 0 and 1.
2) Evaluate fitness for all members and store in **FitnessList[μ]**
3) For **NO_OF_GENERATIONS** generations do
   a) For each member of population $\mathbf{Q_s}$ **[i]**
      i) Produce corresponding offspring $\mathbf{Q_s'}$**[i]** by mutating $\mathbf{Q_s}$ **[i]**
      ii) Evaluate Fitness for all $\mathbf{Q_s'}$**[i]** and store in **OffspringFitnessList[μ]**
   b) Select the best μ members from $\mathbf{Q_s'}$**[μ]** and $\mathbf{Q_s}$ **[μ]** and replace worse members in $\mathbf{Q_s}$ **[μ]** with new better members from $\mathbf{Q_s'}$**[μ]** (Prefer member from Offspring population is the fitness is equal). Copy corresponding fitness values from **OffspringFitnessList[μ]** to **FitnessList[μ].**
   c) Log best fitness value in population $\mathbf{Q_s}$ **[μ]** in fitness log file. Also log Q_s value corresponding to best fitness value in population. Log η values for all successful mutations.

Figure 2. Pseudo-code for our Evolutionary Algorithm for optimising Q_s value for a given market

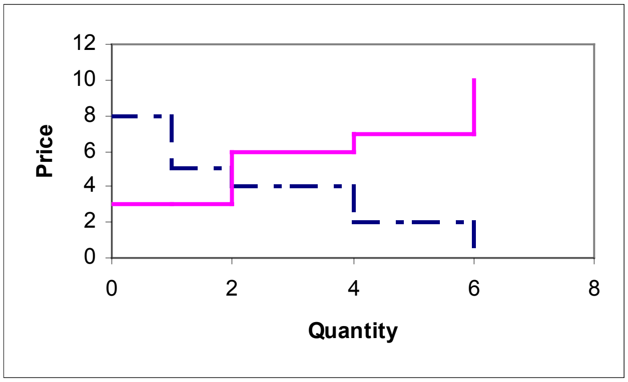

Figure 3 The supply and demand curves for MZIC-I. In this market Po=4 and Qo=2

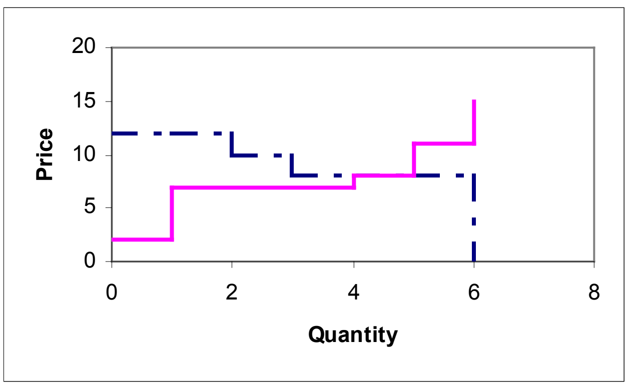

Figure 4 The supply and demand curves for MZIC-II. In this market Po=8 and Qo=5

3.2 Results

3.2.1 MZIC-I

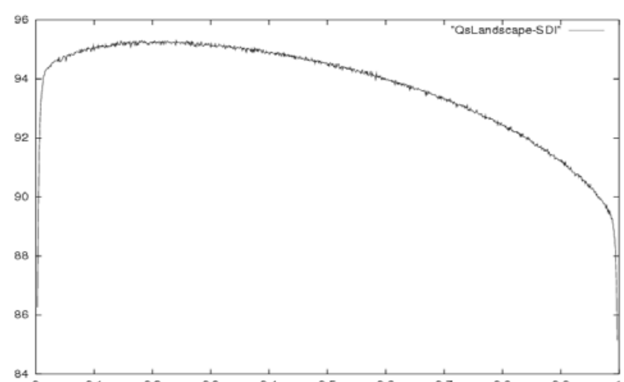

Figure 5 The fitness landscape for the MZIC-I problem. Note that this landscape has been rescaled and plotted for fitness values above 85%. The fitness values are averaged over 10,000 repeats.

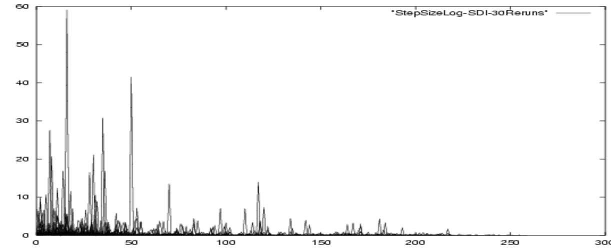

Even by visual inspection one can observe that the MZIC-I landscape is strongly asymmetric and the optimum lies at none of the conventional values of Q_s equal to 1, 0 or 0.5. The step size (η) log shown in confirms the predicted behaviour of the self-adaptive ES as the step sizes for successful offspring get smaller and smaller as the ES hones into the exact optimum. Note that the step size η does not give the value of the actual mutation. The actual mutation value is given by $\eta * N(0,1)$, where $N(0,1)$ is a normally distributed random number with mean 0 and standard deviation 1.

Figure 6 This figure shows the log of the Step Size (η) values for successful offspring as evolution proceeds for all 25 runs of the ES on test case MZIC-I.

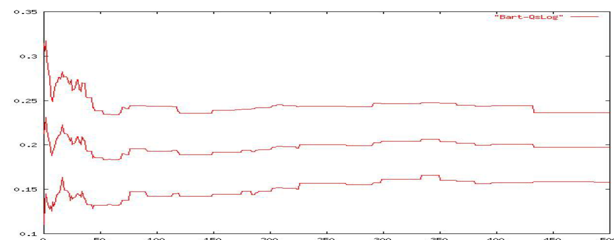

Figure 7 The log of elite Q_s values averaged over 25 runs for MZIC-I. The lines for plus and minus one standard deviation are also shown.

The log for elite Q_s values shown in Figure 7 shows the optimum to lies around Q_s =0.2 (This log depicts information averaged over 25 runs). This agrees with the information depicted in the landscape for the problem, which was shown in Figure 5.

3.2.2 MZIC-II

Figure 8 The fitness landscape for the MZIC-II problem. This landscape is conspicuous because of its seemingly linear nature with most points having 100% fitness. The fitness values are averaged over 10,000 repeats.

The landscape for MZIC-II shown in Figure 8 immediately stands out because of its seemingly linear nature and because of the fact that most of the points in the landscape seem to have 100% efficiency. The elite Q_s values for all 25 runs, which are shown in Figure 9, is extremely noisy and shows poor convergence due to this highly flat landscape, which behaves like a plateau over which the ES displays an unstable behaviour. However, note that the values are clustered around Q_s =0.2.

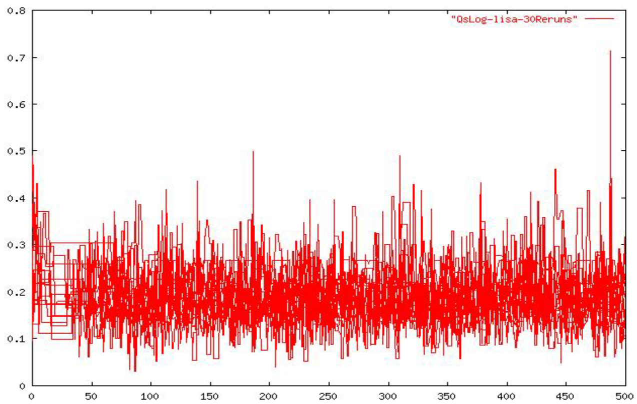

Figure 9 This figure shows the elite Q_s value log for all 25 runs of the ES

A qualitative explanation of why the base-line efficiency in this market is so high can be given as follows. Consider the supply and demand schedule for the MZIC-II problem shown in Figure 4. The only extra-marginal trader in this market who can reduce efficiency by executing a trade is the seller with reserve price 11. Furthermore, the only buyer he can trade with has a reserve price of 12. Given that this buyer always shouts prices below 12 (distributed uniformly in the [0,12] interval) and the extra-marginal seller always shouts prices above 11 – it is highly unlikely that they will ever trade. Hence, 100% efficiency is achieved in most runs. A full quantitative analysis for calculating this likelihood is hard due the presence of the market queue, as well as the fact that we must consider the actions taken by the other traders. But the qualitative explanation gives us a fair understanding of the reason behind the nature of the landscape.

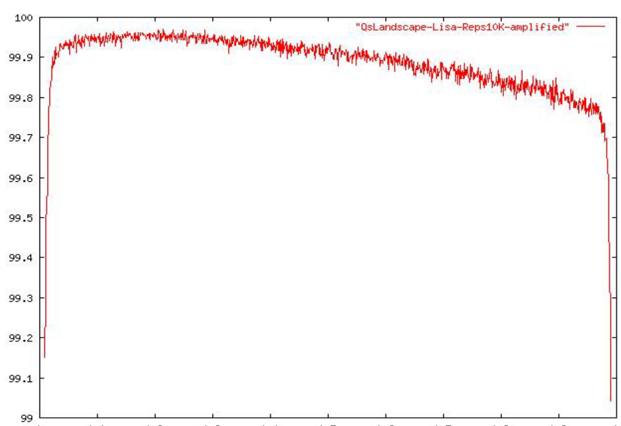

Figure 10 The figure shows the landscape for MZIC-II that has been rescaled and plotted for fitness values above 99%

For cases like MZIC-II, when we amplify the fitness landscape, depicted in Figure 8, it is reassuring to find that the underlying trend does seem to peak at around Q_s =0.2 and confirms the result given by the ES.[1]

4. GAINING INSIGHT IN TO EMPIRICAL RESULTS

The methods we have described in the previous section can help us determine the optimum value of Q_s for a given market, which can then be used to conduct an auction which offers a high likelihood of making larger gains than previously known auction types like the English, Dutch and Continuous Double Auction. However, our technique requires: -

1. A perfect or good knowledge about the nature of the underlying Supply and Demand curves for the market.
2. Significantly large periods of time (of the order of hours) to conduct an evolutionary search through the space of possible market types

An argument against the applicability of such a technique is that the underlying supply and demand data for any given market is unknown and moreover, changes dynamically with time. This is a fair criticism and is perfectly applicable to volatile markets in which the underlying supply and demand curves change rapidly. However, there exist markets which are relatively stable and whose nature changes only slowly with time. In these markets it is possible to determine values of Q_s that can be used for protracted periods of time.

But it would be best if we could determine the optimum Q_s value in spite of imperfect supply and demand data and in a relatively short period of time. Evolutionary techniques are probably ill suited for dynamically determining the optimum Q_s value in a 'live' market. So in order to be able to determine the optimum Q_s in a shorter period of time we need to gain more insight in to the relationship between the optimum Q_s value and the underlying supply and demand curves.

4.1 Generating Test markets

We have decided to parameterise markets in terms of the slope of the supply and demand curves that describe them. We assume that the supply and demand curves are linear in nature. Moreover since the number of traders in our experiments is small (10-22) we must devise a way to define a discrete analogous of supply slope and demand slope. Our markets are defined by a set of six parameters:-

i. **minPrice:** minimum price in the market
ii. **maxPrice:** maximum price in the market
iii. **mS:** the slope of the supply curve
iv. **mB:** the slope of demand curve

[1] Please refer to [5,15] for similar results on the original test-cases M1-M3 used by Cliff [3] as well as other new test cases. There we also provide a more detailed description of the behaviour of the ES in all these test market.

v. **nB:** number of buyers in the market

vi. **nS:** number of sellers in the market

The Supply Curve i.e. the limit prices for sellers are generated as follows: -

- The first seller s_0 is assigned price minPrice.
- Every subsequent trader s_i is allocated a price (minPrice+i*mS) until (minPrice+i*mS)>maxPrice.
    - If (minPrice+i*mS)>maxPrice the limit price for seller s_i is fixed at maxPrice.

The process for generating the demand curve is similar (except the slope mB is interpreted as a negative gradient).

4.2 Gaining Computational leverage

If we vary the slopes through 21 points ([0,20]) each of mB and mS, we have 441 markets in which to need to find the optimum Q_s value. If we set NO_OF_GENERATIONS to 500 and run the ES to completion for each of these markets, this experiment would take more than 3 months to perform on a single machine. To gain computational leverage in this large landscaping process we have made some modifications to our ES. These modifications are based on the observations that very often the ES can stagnate much earlier than the maximum number of generations. The basic idea to is to monitor the behaviour of the ES as it optimises the Q_s value for any market (mB,mS) and terminate its execution if the ES stagnates. We define a parameter MAX_STAGNANT and terminate the ES has been stagnant for the last MAX_STAGNANT generations. There are two criterions for judging if the ES has stagnated.

1. Elite Q_s value has stagnated

We can terminate the execution of the ES if the elite Q_s value has stagnated as this indicates a likely convergence to the optimum Q_s value. We have included a bias in our ES that leads to an offspring replacing a member of the parent population if the fitness of both is equal. So if the elite Q_s value has remained constant – this is probably on account of domination of this value over others rather than a random walk by the ES (which can happen in case of flat (or nearly flat) landscapes). Hence, in this case we terminate the ES and report the elite Q_s value.

2. Elite Fitness value has stagnated

The case in which the fitness value stagnates, but the elite Q_s value keeps changing is trickier. The reader should satisfy herself that there are landscapes in which the efficiency will always be 100% (for eg. in markets in which the supply and demand curves do not intersect). In such cases we have adopted the policy of reporting the average Q_s value – which on an average can be expected to be ~0.5.

4.3 Results

The mB and mS values for this experiment are varied over the interval [0,20] with an increment of 1. The number of buyers and sellers, nB=nS=5. Fitness values are averaged over 1,000 REPEAT_TRIALS. MAX_STAGNANT has a value of 10 for this experiment.

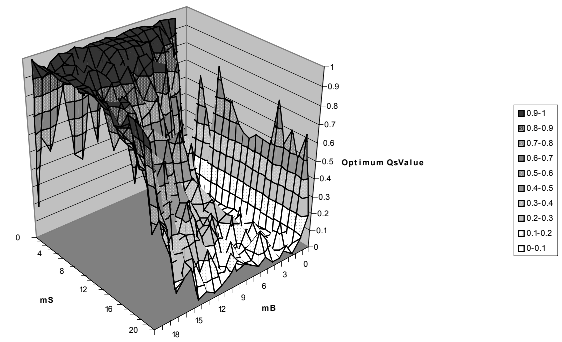

Figure 11 A 3-Dimensional landscape of optimum- Q_s values over a large 2-dimensional space of test markets characterised by the slope of the Supply curve (mS) and Demand curve (mB). See text for discussion

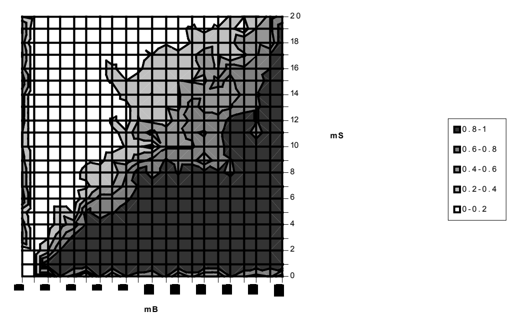

Figure 12 A contour plot of the 3-dimensional landscape shown in Figure 11. See text for discussion

We observe from the following regularities from Figure 12: -

1. Along the axis on which mB=20 almost all markets have an optimum Q_s value of ~1.
2. Conversely when the slope of the supply slope mS=20 most markets have an optimum Q_s value of ~0.
3. A significant number of markets have non-standard optimum Q_s values.
4. We observe that along the mS=mB diagonal in this matrix, there is a consistent trend that which shows the domination of hybrid market types. Interestingly this non-standard (hybrid) Q_s region ($Q_s \in$[0.2,0.4], [0.4,0.6] and [0.6,0.8]) seems to broaden above the mS=20-mB diagonal in this matrix of markets.

We now give some qualitative arguments to show that these observed regularities do agree with intuition. Note that repeated experiment results reported elsewhere confirm these trends. For more results and discussion see [15].

4.3.1 *Along the axis on which mB=20 almost all markets have an optimum Q_s value of ~1*

In cases where the buyer slope mB is high, such that there is only one intra-marginal buyer but many intra-marginal sellers in the market; ensuring 100% efficiency is a question of making sure that the buyer with highest reserve (utility) price trades with the seller with the lowest cost price is, thus generating the largest possible total surplus.

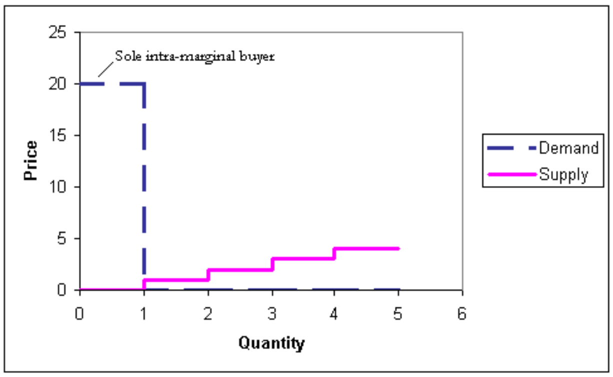

Figure 13 This figure shows the supply and demand schedule for a test market in which mB=20 and mS=2. Note that since there is only one intra-marginal buyer, ensuring 100% efficiency is a question of making sure that she trades with the seller with the lowest cost price.

How does a Q_s value of 1 help us in achieving this aim? A Q_s value of 1 means that only sellers shout offers in the market as the buyers observe silently. This in turn means that sellers are being forced to compete amongst themselves and undercut each other continuously in an effort to steal the only trade that is available in the market. This competitive process leads to the elimination of the efficiency-reducing extra-marginal traders one-by-one as the current-best offer in the market drops below each of their reserve (cost) prices. This process continues till only the seller with the lowest reserve price is left in the market. This seller can now trade with the sole intra-marginal buyer generating 100% efficiency. So in this case a Q_s value of 1 can be seen as an attempt by the mechanism to ensure that only the seller with the lowest cost price trades, thus leading to 100% efficiency. Similar arguments can be made for why a Q_s value of 0 is optimal for markets in which the supply slope mS is high.

4.3.2 *The non-standard (hybrid) Q_s region ($Q_s \in [0.2,0.4]$, $[0.4,0.6]$ and $[0.6,0.8]$) seems to broaden above the mS=20-mB diagonal in this matrix of markets.*

Examining the contour plot shown in Figure 14, we observe that hybrid values of Q_s dominate more and more as the value of the Equilibrium Quantity Q_o decreases. This is shown by the broadening of the hybrid 'bulge' that can be seen along the mS=mB diagonal.

As the value of Q_o decreases, it becomes increasingly critical to get the only profitable trade available in the market 'right'. Since the efficiency achieved is directly related to the surplus generated in this *only* transaction, the effect of using Q_s as a control parameter is heightened. If we look at Figure 14 again in this light we observe that the indeed, the 'hybrid' region does seem to broaden as Q_o decreases. Refer to [14] for details.

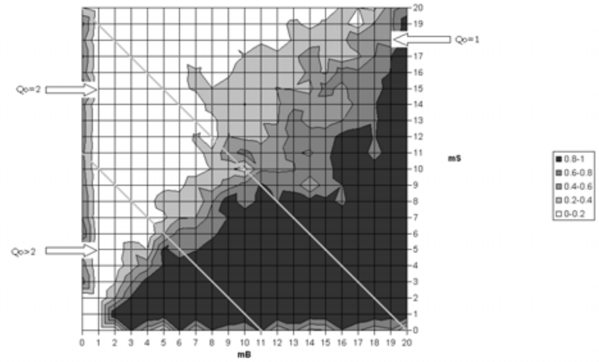

Figure 14 Contour plot with the regions with Equilibrium Quantity Q_o values of 1,2 and greater than 2 marked. See text for discussion.

5. Conclusion and Directions for Future Research

The CDA, even though hugely popular, is also little understood, and gaining insight into its properties is important both from an economic as well as scientific perspective. CDAs are not only the most common type of auction mechanism in the world but they also find application in many artificial intelligence problems using the paradigm of Market Based Control (MBC) [16]. Recent research has shown that the CDA, although difficult to analyse rigorously, lends itself to an experiment-based empirical approach. Our work can be seen as another step in this direction.

Cliff had proposed an idea about market design in which the bidding process could be regulated by a parameter Q_s – our work is an attempt to investigate this idea thoroughly and determine whether it has potential for wide application in actual marketplaces. Our results show that hybrid variants of the CDA are superior to conventional mechanisms for many of our test cases. These are the first-ever results that demonstrate that hybrid markets can dominate conventional variants with more than one type of trading algorithm. More research is needed to prove that these results hold for *all* types of trading strategies and perhaps even human traders. More research is also needed to help us determine exactly how this parameter leads to gains in efficiency in different market contexts. We note that outcomes for experiments performed for the same markets but with different trading strategies can be different. It appears that the optimal Q_s value for a given market is dependent upon the trading strategy used. Experiments with large markets populated with a large number of heterogeneous trading-agents could help us show if this indeed the case and whether hybrid variants emerge as optimal even in these markets.

We recognise that the full potential of this idea for automatic market design can only be realised if we gain better understanding of how it works and are able to devise faster and more accurate ways of designing optimal markets with minimal risk in situations where we have only incomplete or inaccurate information. We have tried to do this by empirically studying patterns which can emerge from conducting a large number of tests over a search space of test-markets characterised by given parameters. We have shown that certain intuitions are confirmed by our experiments and the results agree with some previously known facts about the design of conventional market-mechanisms. We discover that choosing a value of Q_s for a given market can be viewed as an attempt to determine to which extent buyers and sellers (as groups) take on the roles of price 'makers' and price 'takers'. This work shows that the Evolutionary Algorithms are not just optimisation techniques but can be used successfully for search, exploration and knowledge discovery, especially in applications where rigorous quantitative analysis is intractable and the size of the search space is too large for exploration using conventional techniques.

Analysis of multi-player games like markets is hard. We need more insight into the relationship of the optimal Q_s value with underlying market characteristics. Such research can help us in applying these ideas to real-time 'live' markets and ultimately lead to better market dynamics that can help maximise social welfare outcomes.

6. ACKNOWLEDGEMENTS

During his MSc Vibhu Walia has been supported in part by European Commission grant ASI/B7-301/97/0126-08 and in part by a scholarship from the School of Computer Science at the University of Birmingham. Additional support for this project came from the HP Labs Biologically-Inspired Complex Adaptive Systems (BICAS) research group:

http://www.hpl.hp.com/research/bicas

7. REFERENCES

[1] Back T. and Schewefel H. P. " An overview of evolutionary algorithms for parameter optimisation", *Evolutionary Computation*, Vol. 1, No. 1, p. 1-23, 1993

[2] Cliff D. "Evolutionary Optimisation of parameters sets for adaptive software-agent traders in continuous double-auction markets", Technical Report HPL-2001-99, Hewlett-Packard Laboratories, 2001

[3] Cliff D. "Evolution of market mechanism through a continuous space of auction types", Technical Report HPL-2001-326, HP Labs, 2001a

[4] Cliff D. "Evolution of market mechanism through a continuous space of auction-types II: Two-sided auction mechanisms evolve in response to market shocks", Technical Report HPL-2002-128, HP Labs, 2002

[5] Cliff D., Walia V. and Byde A. "Evolved Hybrid Auction Mechanisms in Non-ZIP Trader Marketplaces", Technical Report HPL-2002-247, Hewlett-Packard Laboratories, 2002

[6] Das R., Hanson J. E., Kephart J. O., Tesauro G. "Agent Human Interaction in a Continuous Double Auction", In *Proceedings of the International Joint Conferences on Artificial Intelligence*. Seattle, USA, August 2001

[7] Fogel D. B. *Evolving Artificial Intelligence*, PhD thesis, University of California, San Diego, CA, 1992

[8] Fogel D. B. *Evolutionary Computation: Towards a New Philosophy of Machine Intelligence*, New York, NY: IEEE Press, 1995.

[9] Gode D.K. and Sunder S. "Lower Bounds for Efficiency of Surplus Extraction in Double Auctions", In *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust Eds, Addison Wesley, 1993.

[10] Phelps S., McBurney P., Parsons S., Sklar E. "Co-evolutionary auction mechanism design: a Preliminary Report", Technical Report ULCS-02-019, Department of Computer Science, University of Liverpool, 2002

[11] Preist C., van Tol, M. "Adaptive agents in a persistent double shout auction". *In Proceedings of ICE-98.* ACM 1998

[12] Smith V., "Experimental study of competitive market behavior". *Journal of Political Economy* 70:111-137, 1962

[13] Schwefel, Hans Paul *Numerical Optimisation of Computer Models*, John Wiley and Sons, Chichester, 1981

[14] Schwefel, Hans Paul *Evolution and Optimum Seeking*, John Wiley and Sons, New York, 1995

[15] Walia, V., *Evolving Market Design*, MSc Thesis, School of Computer Science, University of Birmingham, September 2002.

[16] Wellman, Michael P. "The economic approach to artificial intelligence" *ACM Computing Surveys*, 27(3), 340-342, September 1995.

[17] Yao X., Liu Y. and Lin G. "Evolutionary programming made faster," *IEEE Transactions on Evolutionary Computation*, 3(2):82-102, July 1999