Evolved Hybrid Auction Mechanisms in Non-ZIP Trader Marketplaces

Dave Cliff, Vibhu Walia, Andrew Byde
Information Infrastructure Laboratory
HP Laboratories Bristol
HPL-2002-247
September 16th , 2002*

E-mail: dave_cliff@hp.com, vibhu_walia@hp.com, andrew_byde@hp.com

automated mechanism design, market mechanism, market design, ZIP traders, ZI-C traders, genetic algorithms, e-Marketplaces

A previous CIFEr paper demonstrated that a genetic algorithm could be used to automatically discover new optimal auction mechanisms for automated electronic marketplaces populated by software-agent traders. Significantly, the new auction mechanisms are often unlike traditional mechanisms designed by humans for human traders; rather, they are peculiar hybrid mixtures of established styles of mechanism. This previous work concentrated on auction marketplaces populated by software agents running the ZIP trader algorithm (recently shown to outperform human traders). In this paper we provide the first demonstration that qualitatively similar results (i.e., non-standard hybrid mechanism designs being optimal) are also given when similar experiments are performed using a different trader algorithm, namely Gode & Sunder's ZI-C traders. Thus, this paper is the first to show that the previous results were not specific to ZIP traders, and hence it offers significant evidence that evolved hybrid auction mechanisms may be found to out-perform traditional mechanisms for any style of trader-agent.

* Internal Accession Date Only
Approved for External Publication
© Copyright Hewlett-Packard Company 2002

Evolved Hybrid Auction Mechanisms in Non-ZIP Trader Marketplaces

Dave Cliff [1], Vibhu Walia[1,2], Andrew Byde[1]
[1:]Hewlett-Packard Laboratories, Filton Road, Bristol BS34 8QZ, England, U.K.
[2:]School of Computer Science, University of Birmingham, Birmingham B15 2TT, England, U.K.

dave_cliff@hp.com vibhu_walia@hp.com andrew_byde@hp.com

Abstract: A previous CIFEr paper demonstrated that a genetic algorithm could be used to automatically discover new optimal auction mechanisms for automated electronic marketplaces populated by software-agent traders. Significantly, the new auction mechanisms are often unlike traditional mechanisms designed by humans for human traders; rather, they are peculiar hybrid mixtures of established styles of mechanism. This previous work concentrated on auction marketplaces populated by software agents running the ZIP trader algorithm (recently shown to outperform human traders). In this paper we provide the first demonstration that qualitatively similar results (i.e., non-standard hybrid mechanism designs being optimal) are also given when similar experiments are performed using a different trader algorithm, namely Gode & Sunder's ZI-C traders. Thus, this paper is the first to show that the previous results were not specific to ZIP traders, and hence it offers significant evidence that evolved hybrid auction mechanisms may be found to outperform traditional mechanisms for *any* style of trader-agent.

Submitted in abridged form to CIFEr'03.

I. INTRODUCTION

ZIP (Zero-Intelligence-Plus) artificial trading agents, introduced in 1997 [1], are software agents (or "robots") that use simple machine learning techniques to adapt to operating as buyers or sellers in open-outcry auction-market environments similar to those used in the experimental economics work of Smith (e.g. [2]). Although initially developed purely to address deficiencies in Gode & Sunder's ZI-C traders [2], recent experimental work by Das *et al.* at IBM [4] has shown that ZIP traders (unlike ZI-Cs) consistently out-perform human traders in human-against-robot auction marketplaces.

The operation of ZIP traders has been successfully demonstrated in experimental versions of continuous double auction (CDA) markets similar to those found in the international markets for commodities, equities, capital, and derivatives; and in posted-offer auction markets similar to those seen in domestic high-street retail outlets [1,2]. In any such market, there are a number of parameters that govern the adaptation and trading processes of the ZIP traders. In the original formulation [1], the values of these parameters were set by hand, using "educated guesses". However, at CIFEr'98, the first results were presented from using a standard genetic algorithm (GA) to automatically optimise these parameter values [5], thereby eliminating the need for skilled human input in deciding the values of the parameters; more details of these GA results were subsequently given in [6].

In all previous work using artificial traders, ZIP or otherwise, the market mechanism (i.e., the type of auction the traders are interacting within) had been fixed in advance. Well-known market mechanisms from human economic affairs include: the English auction (where sellers stay silent and buyers quote increasing bid-prices), the Dutch Flower Auction (where buyers stay silent and sellers quote decreasing offer-prices); the Vickery or second-price sealed-bid auction (where sealed bids are submitted by buyers, and the highest bidder is allowed to buy, but at the price of the *second-highest* bid -- this curious mechanism encourages honesty and is robust to attack by dishonest means); and the CDA (where sellers announce decreasing offer prices while *simultaneously and asynchronously* the buyers announce increasing bid prices, with the sellers being free to accept any buyer's bid at any time and the buyers being free to accept any seller's offer at any time).

At CIFEr'02, Cliff [7] presented the first ever results from experiments where a GA optimised not only the parameter values for the trading agents, but also the style of market mechanism in which the traders operate. To do this, a space of possible market mechanisms was created for evolutionary exploration. The space included the CDA and also one-sided auctions similar (but not actually identical to) the English Auction (EA) and the Dutch Flower Auction (DFA); and significantly this space is *continuously variable*, allowing for any of an *infinite* number of peculiar hybrids of these auction types to be evolved, which have no known correlate in naturally occurring market mechanisms. While there was nothing to prevent the GA from settling on solutions that correspond to the known CDA auction type or the EA-like and DFA-like one-sided mechanisms, Cliff [7,8] repeatedly found that the GA settles on hybrid solutions and that these hybrids lead to the most desirable market dynamics. Although the hybrid market mechanisms could easily be implemented in online electronic marketplaces, they have not been designed by humans: rather they are the product of evolutionary search through a continuous space of possible auction-types. Thus, the CIFEr'02 paper [7] was the first ever demonstration that radically new market mechanisms for artificial traders may be designed by automatic means, thereby establishing the new field of *automated mechanism design.* Independently, some similar work was under development elsewhere, and was published a couple of months later [9].

As all of Cliff's results [7,8] were from marketplaces populated by ZIP traders, an obvious question to ask is to what extent those results were dependent on the use of ZIP traders. That is: if non-ZIP trader-agents had been used, would similar hybrid auction mechanisms still be found to be optimal by the GA? In this paper we present the first demonstration that

Cliff's [7,8] results were not specific to the use of ZIP traders. Section II gives an overview of the background experimental methods and results as published by Cliff [5,6,7,8]. In Section III we present new ZIP-trader data, visualising the "fitness landscapes" explored by the GA. In Section IV we then show comparable fitness landscapes calculated for Gode & Sunder's ZI-C trader-agents. In Section V we discuss the strong qualitative similarities between the ZIP and ZI-C landscapes, and conclude that they offer firm evidence for the claim that automated mechanism design can find hybrid auction styles that give desirable performance, whatever the style of trader-agent used in the marketplace. Note that, in the rest of this paper, we use $v=\boldsymbol{U}[x,y]$ to denote a random real value v generated from a uniform distribution over the range $[x,y]$.

II. BACKGROUND

A. Zero-Intelligence Plus (ZIP) Traders

ZIP traders are described fully in [1], which includes sample source-code in the *C* programming language. For the purposes of this paper a high-level description of the key parameters is sufficient. Each ZIP trader i is given a private (secret) limit-price, λ_i, which for a seller is the price below which it must not sell and for a buyer is the price above which it must not buy. If a ZIP trader completes a transaction at its λ_i price then it generates zero utility ("profit" for the sellers or "saving" for the buyers). For this reason, each ZIP trader i maintains a time-varying margin $\mu_i(t)$ and generates quote-prices $p_i(t)$ at time t according to $p_i(t)=\lambda_i\ (1+\mu_i(t))$ for sellers and $p_i(t)=\lambda_i\ (1-\mu_i(t))$ for buyers. The "aim" of traders is to maximise their utility over all trades, where utility is the difference between the accepted quote-price and the trader's λ_i value. Trader i is given an initial value $\mu_i(0)$ (i.e., $\mu_i(t)$ for $t=0$) which is subsequently adapted over time using a simple machine learning technique known as *the Widrow-Hoff rule* which is also used in back-propagation neural networks. This rule has a "learning rate" parameter β_i that governs the speed of convergence between trader i's quoted price $p_i(t)$ and the trader's idealised "target" price $\tau_i(t)$. When calculating $\tau_i(t)$, traders introduce a small random absolute perturbation generated from $\boldsymbol{U}[0,c_a]$, and also a small random relative perturbation coefficient generated from $\boldsymbol{U}[1-c_r,1]$ (when a trader is reducing its $p_i(t)$) or $\boldsymbol{U}[1,1+c_r]$ (when increasing $p_i(t)$) where c_a and c_r are global system constants. To smooth over noise in the learning system, there is an additional "momentum" parameter γ_i for each trader (such momentum terms are also commonly used in back-propagation neural networks).

Thus, adaptation in each ZIP trader i has the following parameters: initial margin $\mu_i(0)$; learning rate β_i; and momentum term γ_i. In an entire market populated by ZIP traders, these three parameters are assigned to each trader from uniform random distributions each of which is defined via "min" and "delta" values in the following fashion: $\mu_i(0)=\boldsymbol{U}(\mu_{min}, \mu_{min}+\mu_\Delta)$; $\beta_i=\boldsymbol{U}(\beta_{min}, \beta_{min}+\beta_\Delta)$; and $\gamma_i=\boldsymbol{U}(\gamma_{min}, \gamma_{min}+\gamma_\Delta)$.

Hence, to initialise an entire ZIP-trader market it is necessary to specify values for the six market-initialisation parameters $\mu_{min}, \mu_\Delta, \beta_{min}, \beta_\Delta, \gamma_{min}$, and γ_Δ; and also for the two system constants c_a and c_r. And so it can be seen that any set of initialisation parameters for a ZIP-trader market exists within an eight-dimensional real space, conventionally denoted by $\mathbf{R}^8$. Vectors in this 8-space can be considered as genotypes, and from an initial population of such genotypes it is possible to allow a GA to find new genotypes that best satisfy an appropriate evaluation function. This is exactly the process that was introduced at CIFEr'98 [5,6], as described in Section II.C below. Before that, we discuss the issue of simulating the passage of time.

When monitoring events in a real auction, as more precision is used to record the time of events, so the likelihood of any two events occurring at exactly the same time is diminished. For example, if two quotes made at five minutes past nine are both recorded as occurring at 09:05, then they appear in the record as simultaneous; but a more accurate clock would have been able to reveal that the first was made at 09:05:01.64 and the second at 09:05:01.98. Even if two events occur absolutely at the same time, very often some random process (e.g. what direction the auctioneer is looking in) acts to break the simultaneity.

Thus, we may simulate real marketplaces (and implement electronic marketplaces) using techniques where each significant event always occurs at a unique time. We may choose to represent these by real high-precision times, or we may abstract away from precise time-keeping by dividing time into discrete (possibly irregular) *slices*, numbered sequentially, where one significant event is known to occur in each slice. Such a time-slicing approach was used in previous work [1,5,6,7,8]. In each time-slice, the atomic "significant event" is one quote being issued by one trader and the other traders then responding either by ignoring the quote or by one of the traders accepting the quote. (NB in [4] a continuous-time formulation of the ZIP-trader algorithm was used).

In the markets described here and in [1,5,6], on each time-slice a ZIP trader i is chosen at random from those currently able to quote (i.e. those who hold appropriate stock or currency), and trader i's quote price $p_i(t)$ then becomes the "current quote" $q(t)$ for time t. Next, all traders j on the contraside (i.e. all buyers j if i is a seller, or all sellers j if i is a buyer) compare $q(t)$ to their own current quote price $p_j(t)$ and if the quotes cross (i.e. if $p_j(t)<=q(t)$ for sellers, or if $p_j(t)>=q(t)$ for buyers) then the trader j is able to accept the quote. If more than one trader is able to accept, one is chosen equiprobably at random to make the transaction. If no traders are able to accept, the quote is regarded as "ignored". Once the trade is either accepted or ignored, the traders update their $\mu(t)$ values using the learning algorithm outlined above, and the current time-slice ends. This process repeats for each time-slice in a trading period, with occasional injections of fresh currency and stock, or redistribution of λ_i limit prices, until a maximum number of time-slices have completed.

B. Space of Possible Auctions

Now consider the case where we implement a ZIP-trader continuous double auction (CDA) market. In any one time-slice in a CDA either a buyer or a seller may quote, and in the definition of a CDA a quote is equally likely from each side. One way of implementing a CDA is, at the start of each time-slice, to generate a random binary variable to determine whether the quote will come from a buyer or a seller, and then to randomly choose one individual as the quoter from whichever side the binary value points to. Here, as in previous ZIP work [1,5,6] the random binary variable is always independently and identically distributed over all time-slices.

Let $Q=b$ denote the event that a buyer quotes on any one time-slice and let $Q=s$ denote the event that a seller quotes, then for the CDA we can write $Pr(Q=s)=0.5$ and note that because $Pr(Q=b)=1.0-Pr(Q=s)$ it is only necessary to specify $Pr(Q=s)$, which we will abbreviate to Q_s hereafter. Note additionally that in an EA we have $Q_s=0.0$, and in the DFA we have $Q_s=1.0$. Thus, there are at least three values of Q_s (i.e. *0.0, 0.5,* and *1.0*) that correspond to three types of auction familiar from centuries of human economic affairs.

However, although the ZIP-trader case of $Q_s=0.5$ is indeed a good approximation to the CDA, the fact that any ZIP trader j will accept a quote whenever $q(t)$ and $p_j(t)$ cross means that the one-sided extreme cases $Q_s=0.0$ and $Q_s=1.0$ are not exact analogues of the EA and DFA. Nevertheless, consider the implications of considering values of Q_s of 0.0, 0.5, and 1.0 not as three distinct market mechanisms, but rather as three points on a *continuum*. How do we interpret, for example, $Q_s=0.1$? Certainly there is a straightforward implementation: on the average, for every nine quotes by buyers, there will be one quote from a seller. Yet the history of human economic affairs offers no examples (as far as we are aware) of such markets: why would anyone suggest such a bizarre way of operating, and who would go to the trouble of arbitrating (i.e., acting as an auctioneer for) such a mechanism? Nevertheless, there is no *a priori* reason to argue that the three known points on this Q_s continuum are the only loci of useful auction types. Maybe there are circumstances in which values such as $Q_s=0.1$ are preferred. Given the infinite nature of a real continuum, it seems appealing to use an automatic exploration process, such as a GA, to identify useful Q_s values.

Thus, a ninth dimension was added to the search space, and the genotype in the GA is now the eight real values for ZIP-trader initialisation, plus a real value for Q_s, so the GA is searching for points in $\mathbf{R}^9$ that give the best market dynamics.

C. The Genetic Algorithm

A simple genetic algorithm was used. In each experiment reported in [5,6,7,8] a population of size 30 was used, and evolution was allowed to progress for some number of generations n_g. In each generation, all individuals were evaluated and assigned a "fitness" value (reflecting how good that genotype's market dynamics were); and the next generation's population was then generated via mutation and crossover on parents identified using rank-based selection. Elitism (where an unadulterated copy of the fittest individual from generation g is inserted into the population of $g+1$) was also used.

The genome of each individual was simply a vector of nine real values. In each experiment, the initial random population was created by generating random values from $\boldsymbol{U}[0,1]$ for each locus on each individual's genotype. Crossover points were between the real values, and crossover was governed by a Poisson random process with an average of between one and two crosses per reproduction event. Mutation was implemented by adding random values from $\boldsymbol{U}[-m(g),+m(g)]$ where $m(g)$ is the mutation limit at generation g (starting the count at $g=0$). Mutation was applied to each locus in each genotype on each individual generated from a reproduction event, but the mutation limit $m(g)$ was gradually reduced via an exponential-decay annealing function of the form: $log_{10}(m(g))=-(log_{10}(m_s)-(g/(N_g-1))log_{10}(m_s/m_e))$ where N_g is the maximum number of generations and m_s is the "start" mutation limit (i.e., for $m(0)$) and m_e is the "end" mutation limit (i.e., for $m(n_g-1)$). In all the experiments reported here and in earlier papers [7,8], $N_g=10^3$, $m_s=0.05$, and $m_e=0.0005$.

If ever mutation caused the value at a locus to fall outside *[0.0,1.0]* it was simply clipped to stay within that range. This clip-to-fit approach to dealing with out-of-range mutations biases evolution toward extreme values (i.e. the upper and lower bounds of the clipping), and so Q_s values of 0.0 or 1.0 are, if anything, more likely than values within those bounds. Initial and mutated genome values of μ_Δ, β_Δ, and γ_Δ were also clipped to keep genome vectors within the unit hypercube, i.e. to satisfy the constraints $(\mu_{min}+\mu_\Delta)<=1.0$, $(\beta_{min}+\beta_\Delta)<1.0$, and $(\gamma_{min}+\gamma_\Delta)<1.0$.

The fitness of genotypes was evaluated using the same methods as described in [5,6,7,8]. One *trial* of a particular genome was performed by initialising a ZIP-trader market from the genome, and then allowing the ZIP traders to operate within the market for a fixed number of trading periods, with allocations of stock and currency being replenished between trading periods. Each trading period ended either when no more trades are possible, or a maximum number of time-slices is reached.

During each trading period, Smith's α measure [2] of deviation of transaction prices from the theoretical market equilibrium price was monitored, and a front-weighted average was calculated across the trading periods in the trial. As the outcome of any one such trial is influenced by stochasticity in the system, the final fitness value for an individual was calculated as the arithmetic mean of 100 such trials. Note that as minimal deviation of transaction prices from the theoretical equilibrium price is desirable, lower scores are better: the intention here is to *minimise* the fitness value.

In the CIFEr'02 paper [7] the number of generations n_g for each experiment was set to equal N_g (i.e., 1000), but all the significant evolutionary activity was found to occur in the first 500 generations; hence in subsequent work [8] n_g=500 was used. Thus, in any one experiment, there are 30 individuals evaluated over 500 generations where each evaluation involves calculating the mean of 100 trials, so a total of 1.5 million market trials would be executed in any one GA experiment. Nevertheless, the progress of each GA experiment is itself affected by stochasticity (e.g. the GA may become trapped on local optima) and so to generate reliable results each experiment was repeated 50 times, requiring a total of 75 million market trials. On a current single-CPU PC, 50 repetitions of the single-schedule experiments from [7] take around four days to complete, while 50 repetitions of the dual-schedule experiments from [8] take nearer eight days. New results from 45 such 50-repeat ZIP-trader experiments are shown in Section III, which would have required over 250 days of continuous processing had a single CPU been used.

D. Previous Results

In the CIFEr'02 paper [7], three differing market supply and demand schedules were used, shown here in Figures 1, 2, and 3, and hereafter referred to as markets M1, M2, and M3 respectively. Each of Figures 1 to 3 shows a supply and demand schedule for a marketplace with 11 buyers and 11 sellers, each empowered to buy/sell one unit of commodity, and all three are similar (or identical) to the schedules used by Smith [2]. Figure 4 shows results from 50 repetitions of an experiment where the GA explores the $\mathbf{R}^9$ subspace in an attempt to optimise the ZIP-trader market parameters for operating in M1: for each experiment, the fitness of the best (elite) member of the population is recorded. The results are clearly tri-modal. Of the 50 repetitions, in five the elite ends up on fitness minima of about 3.2, while the other two elite fitness modes are on less-good minima of around 4.0 and 4.75. For comparison, Figure 5 shows the results of 50 repeats of the same experiment, where the value of Q_s was *not* evolved, being instead clamped at 0.5: i.e. the CDA value. The CDA mechanism is often applauded as an auction mechanism in which equilibration is rapid and stable, so we could expect the best fitness from using this market type. With the fixed CDA auction style, an average elite fitness of around 4.5 is settled on by the majority of experiments (48 repetitions) while a small minority (2 repetitions) settle on a less good mode of around 5.1. Clearly then, the evolved-mechanism results are better than the fixed-mechanism CDA results; that is, when the GA is allowed so find its own value of Q_s rather than have the CDA Q_s value of 0.5 imposed on it, it finds fitter solutions – solutions with less deviation of transaction prices from the equilibrium price. As it happens, the Q_s value found in the best elite mode for the evolving-mechanism M1 experiments is zero [7], and for M2 the best Q_s was also zero [7]. But, surprisingly and significantly, for M3 the best Q_s was neither zero, nor 0.5, nor 1.0 – i.e. none of the Q_s values corresponding to traditional human-designed auction mechanisms. Rather, for M3, the best Q_s value was found to be around 0.16 [7].

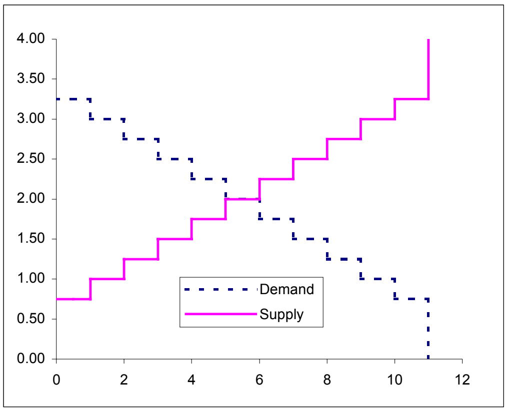

Figure 1: Supply and demand schedules for market M1. Vertical axis is Price; horizontal axis is Quantity.

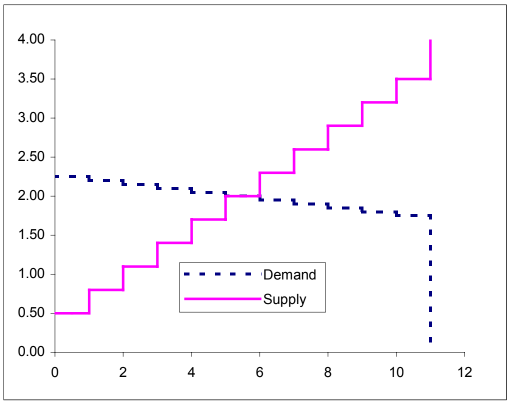

Figure 2: Supply and demand schedules for market M2.

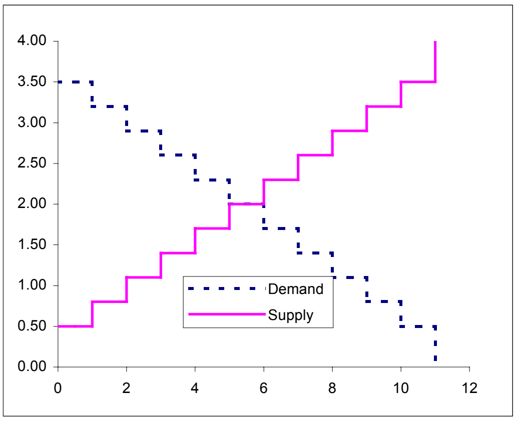

Figure 3: Supply and demand schedules for market M3.

All of the results in the CIFEr'02 paper came from experiments in which the same static supply/demand schedule was used for the duration of each evaluation of every genotype. This is a somewhat unrealistic simplification, for two reasons. First, a primary reason why auction mechanisms such as the CDA are of interest is their ability to adapt to changes in the market's supply and demand curves. Second, it is likely that the GA exploited this regularity and *over-fitted* the ZIP-trader parameters to the particular market schedules used (e.g. a genome that does well in M1may perhaps perform poorly in M2). Thus, in a subsequent paper [8], similar experiments were run but in these new studies the evaluation of a genotype involved six trading periods on one market schedule, followed by a shock-change to another schedule, and then another six trading periods on the new schedule; with the fitness of the genotype being calculated over the entire twelve periods of trading.

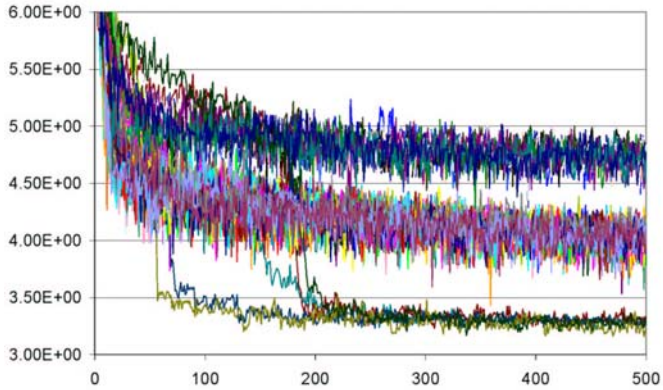

Figure 4: Elite fitness values from 50 repetitions of the 1000-generation evolving-mechanism (EM) experiment operating with M1. Lower values are better solutions (less deviation from equilibrium). Results are trimodal, with five of the repetitions (10%) settling to values around 3.2.

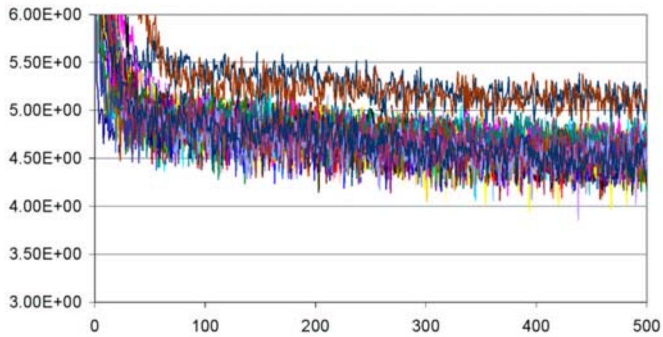

Figure 5: Elite fitness values from 50 repetitions of a 500-generation experiment operating with M1, but with a fixed-mechanism (FM) CDA of Q_s*=0.5*: bimodal results, with 96% of the repetitions settling to fitness values around 4.5 and the remaining 4% at around 5.2.

Hence in these experiments the genotypes had to optimise not only the ZIP-trader's *ab initio* adaptation to the first schedule but also their re-adaptation to the new schedule introduced half-way through the evaluation process. Two sets of experiments were performed: one in which the ZIP traders operated in M1 for six periods followed by a shock-change to M2 for the final six periods (referred to as the M1M2 experiments); and another in which the order was M2 followed by M1 (referred to as M2M1). It was demonstrated [8] that the order was significant: the M1M2 results differed significantly from the M2M1 results. Although in the single-schedule experiments both M1 and M2 were found to have optima at Q_s=0, when the two schedules were both used in one trial then non-zero values of Q_s evolved: for M1M2 the best-mode value was a "hybrid" of around 0.25; while for M2M1 the best value was 0.45, which did not yield statistically significant differences in performance from the CDA value of 0.5.

Having established the background to our current work, we now proceed with introducing our new results.

III. ZIP-TRADER FITNESS LANDSCAPES

A. Methods and Results

As was stated in the previous section, the genotypes in the ZIP-trader experiments are within $\mathbf{R}^9$ (strictly, they are all within the real unit hypercube $[0.0,1.0]^9$). For any such genotype, one evaluation (e.g. taking the mean score from 100 trials, as used here) will give a fitness score for that genotype; and so it is possible in principle to visualise the "fitness landscape" as a surface over the 9-d axes of the genotype space. Visualising such a 10-d object in the two or three dimensions that we humans are familiar with communicating in is manifestly problematic; yet appropriate visualisations can be highly valuable in demonstrating that the results of the GA's evolutionary search are indeed a plausible global optimum. Thus, in this section, we present new data showing visualisations of the fitness landscapes for all five of the ZIP-trader experiments reported in [7] and [8] (i.e., M1, M2, M3, M1M2, and M2M1), before showing in Section IV the fitness landscapes from comparable experiments where the markets are populated by ZI-C traders.

To understand the visualisation, consider Figure 5. In this set of fifty M1 experiments the value of Q_s was fixed at 0.5 and by generation 500 there are two clear elite-fitness modes: one at approx 4.5 and one at approx 5.2. Of the fifty repetitions, 96% settle to the first mode and 4% settle to the second. This could be represented by a histogram where the horizontal axis represents discretized ("binned") values of the elite-fitness mode, and the vertical axis represents the frequency with which each mode is observed; for the M1 Q_s=0.5 data of Figure 5 we would see two distinct peaks in the histogram: a big one around 4.5 and a smaller one around 5.2.

Now to visualise the entire fitness landscape for ZIP traders in M1, run more fixed-mechanism experiments but for each set of 50 repetitions hold the value of Q_s fixed at some value while all the other 8 ZIP parameters on the genome are optimised by the GA. These data allow us to plot a 3-d projection of the 10-d fitness landscape: in our projection, one horizontal dimension is the fixed value of Q_s; another is the elite-fitness mode-value; and the vertical axis shows the frequency with which the different mode values are reached for each of the Q_s values. Figure 6 shows a perspective projection of such a 3-d histogram, calculated for ZIP traders in M1. An alternative view of the same data is presented in Figure 7, i.e. as a contour plot with a logarithmic compression function applied to the frequency values. For comparison, Figures 8, 9, 10, and 11 show such contour plots of the fitness landscape for ZIP-trader experiments with markets M2, M3, M1M2, and M2M1 respectively.

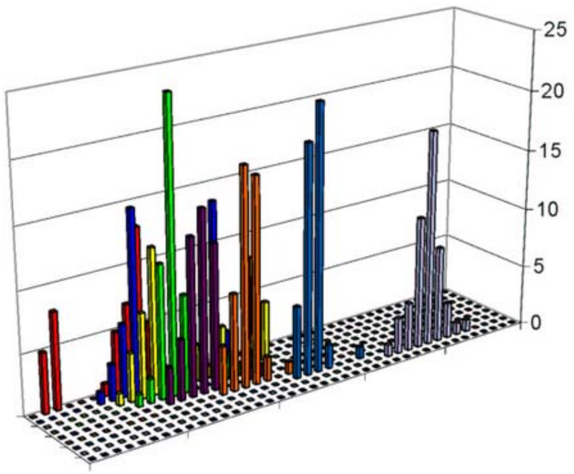

Figure 6: 3-d histogram showing fitness-mode frequency data from multiple repetitions of ZIP-trader GA experiments in M1 over a variety of fixed Q_s values. The narrow horizontal axis is Q_s (from 0.0 at the rear to 1.0 at the front, increasing at intervals of 0.125); the long horizontal axis is elite-fitness value from 3.0 at the left to 8.0 at the right, in "bins" of 0.125; the vertical axis shows the frequency with which the GA settles to each elite-fitness (out of 50 repetitions at each fixed Q_s value). Note that the elite-fitness values for Q_s=1.0 are so poor (i.e., so high) that their histogram data lie off the scale to the right.

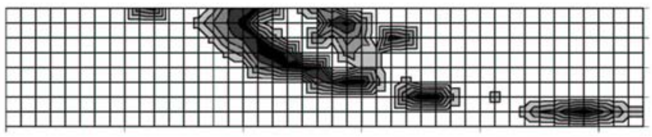

Figure 7: Contour plot of the data shown in Figure 6. Horizontal axis is fitness values from 2.0 at the left to 7.5 at the right (grid-spacing is 0.125); vertical axis is Q_s from 0.0 at the top to 1.0 at the bottom (grid-spacing is 0.125). Darker shading represents higher frequency. Nonzero frequencies for Q_s=1.0 are so poor that they lie off the scale to the right.

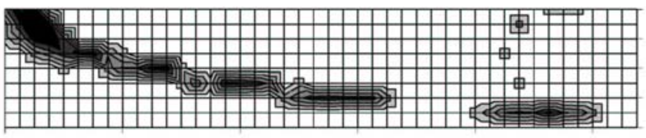

Figure 8: Contour plot of fitness landscape for ZIP traders in M2. Scale as for Fig. 7. Nonzero frequencies for Q_s=1.0 lie off the scale to the right.

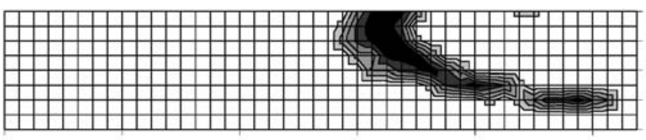

Figure 9: Contour plot of fitness landscape for ZIP traders in M3. Scale as for Figure 7. Fitness values for Q_s>=0.875 lie off the scale to the right.

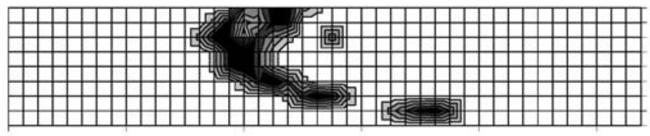

Figure 10: Contour plot of fitness landscape for ZIP traders in M1M2. Scale as for Figure 7. Nonzero frequencies for Q_s=1.0 lie off the scale to the right.

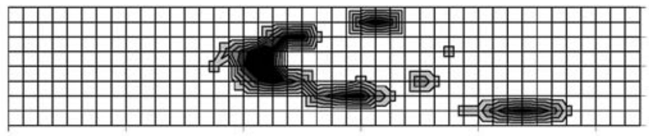

Figure 11: Contour plot of fitness landscape for ZIP traders in M2M1. Scale as for Figure 7. Nonzero frequencies for Q_s=0.0 and Q_s=1.0 lie off the scale to the right.

B. Discussion

All of the contour plots of the fitness landscapes show good agreement with the previous experimental results, in the sense that the minima (i.e. left-most data points) on the contour plots are in good agreement with the values discovered by the evolving-mechanism GA experiments reported on in [7,8] and summarised in Section II.D: markets M1 and M2 both have minima at Q_s=0.0; M3 has a minimum between Q_s=0.125 and Q_s=0.25; M1M2 has a clear minimum at Q_s=0.25; and M2M1 has a minimum at Q_s=0.5. It is also worth noting that all the contour plots show some degree of multi-modality for some values of Q_s.

Thus, the evolving-mechanism results from [7,8] are supported by this brute-force exploration of the fixed-mechanism fitness landscapes for each market schedule: in each case, the evolved value of Q_s is very close to the value identified by empirical examination of the fitness landscapes, and the multi-modality of each fitness landscape justifies the use of multiple repetitions of each experiment in order to identify the true optimal solution.

Having established that for ZIP-traders the fitness landscapes are good illustrations of the optimising performance of the evolving-mechanism GA, we now go on to demonstrate that qualitatively similar results can be observed in marketplaces populated by non-ZIP traders.

IV. FITNESS LANDSCAPES FOR ZI-C TRADERS

A. Methods

To explore whether the previous results were specific to ZIP traders, we ran a new series of parallel experiments, using a different trader algorithm: Gode & Sunder's "zero intelligence constrained" (ZI-C) algorithm [3]. The ZI-C trader algorithm is minimally simplistic: each trader generates a quote price at random from a uniform distribution over $\boldsymbol{U}[q_{min}, \lambda_i]$ for buyers and $\boldsymbol{U}[\lambda_i , q_{max}]$ for sellers, where q_{min} and q_{max} are system constants. Quotes that "cross" lead to a trade. That is: if ever a buyer's bid-price is greater than the last seller's offer price, the offer is considered "accepted"; and if ever a seller's offer-price is less than the last buyer's bid-price, the bid is considered "accepted".

The ZI-C traders are so simple that there are no control parameters for which it would be sensible to employ a GA to optimise the values of. Nevertheless, we can simply study the performance of markets populated by ZI-C traders as we vary the value of Q_s in those markets, deriving fitness landscapes comparable to those shown in Figures 7 to 11.

Gode & Sunder's central result in [3] was the demonstration that CDA markets populated by ZI-C traders could generate allocative efficiency scores close to 100% -- very similar to the values that human traders scored in similar or identical CDA market experiments. As ZI-C traders have some severe difficulties in equilibrating in asymmetric market schedules such as M2 [1], it is not appropriate to use Smith's α measure as the basis of the fitness function. Rather, we simply calculate the allocative efficiency of the ZI-C marketplace at a number of values of Q_s between zero and one. If ZI-C traders perform best when using a "standard" mechanism design, we would expect to see peak performance (i.e., highest values of allocative efficiency) at Q_s values of zero, one, or 0.5. Significantly, these peaks do not always appear. For full details of these ZI-C experiments, see [10].

B. Results

For each market schedule M1, M2, and M3, ZI-C trader allocative efficiencies were calculated by taking the arithmetic mean of 10,000 market trials at values of Q_s ranging from 0.0 to 1.0 in steps of 0.001. In all three cases, there were severe sharp attenuations in allocative efficiencies at the extreme Q_s values of 0.0 and 1.0, caused by the fact that trades only happen when quotes from a buyer and a seller "cross" (as described above) and so if there are no quotes coming from one side of the market, no trades can occur. Allocative efficiency results are shown for ZI-Cs operating in markets M1, M2, and M3 in Figures 12, 13, and 14 respectively. Note that the data in all three figures show clear asymmetries, and for M1 and M3 the peak performance occurs at hybrid values of Q_s different from the standard values of 0.0, 0.5, and 1.0.

As further exploration, such experiments were also run with ZI-C traders operating under two new market schedules not used in the previous work. In the each of these new schedules, referred to here as M4 and M5, there are six sellers and six buyers; each empowered to sell or buy one unit at a given private limit price. In M4 the seller limit prices are 3, 3, 6, 6, 7, and 7; while the buyer limit prices are 8, 5, 4, 4, 2, and 2. In M5 the seller limits are 2, 6, 6, 6, 7, and 8; and the buyer limits are 8, 8, 4, 4, 2, and 2. These supply and demand schedules are illustrated in Figures 15 and 16 respectively. The resulting ZI-C allocative-efficiency landscapes for M4 and M5 are shown in Figures 17 and 18 respectively. Again, these figures show clear asymmetries and have maximal fitness at hybrid Q_s values.

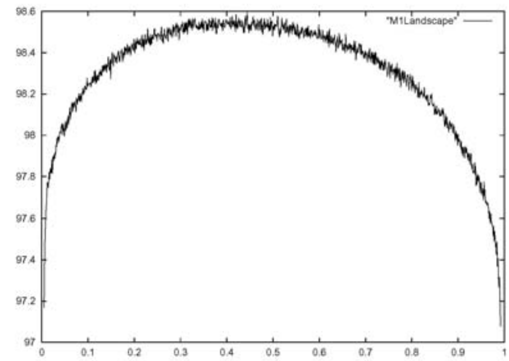

Figure 12: M1 ZI-C fitness landscape, showing allocative efficiency (vertical axis: 97.0 to 98.6) obtained for Q_s=0.0 to Q_s=1.0 (horizontal axis). Search by a GA reveals the peak fitness to be around Q_s=0.39.

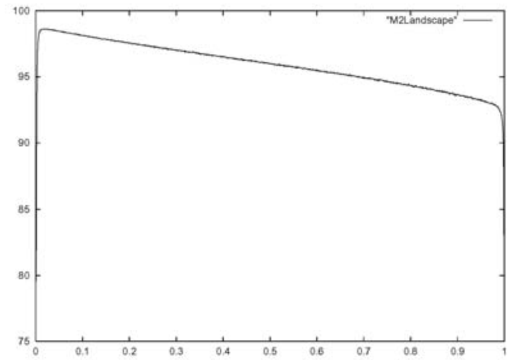

Figure 13: M2 ZI-C fitness landscape, showing allocative efficiency (vertical axis: 75.0 to 100.0) obtained for Q_s=0.0 to Q_s=1.0 (horizontal axis). Search by a GA reveals the peak fitness to be around Q_s=0.01.

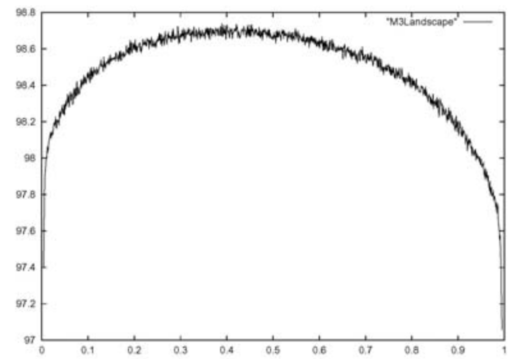

Figure 14: M3 ZI-C fitness landscape, showing allocative efficiency (vertical axis: 97.0 to 98.8) obtained for Q_s=0.0 to Q_s=1.0 (horizontal axis). Search by a GA reveals the peak fitness to be around Q_s=0.42.

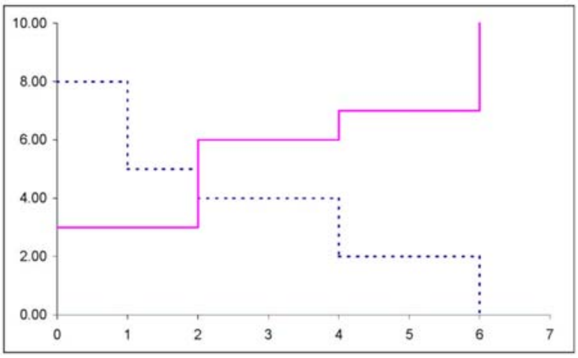

Figure 15: Supply and demand schedules for market M4.

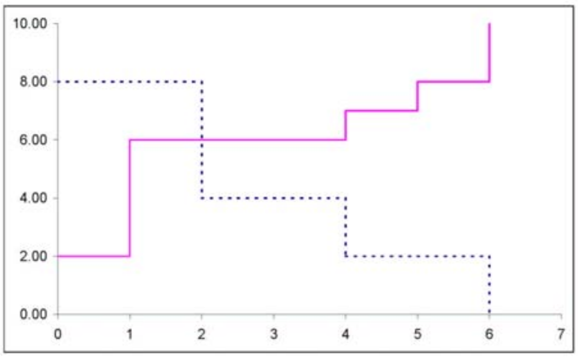

Figure 16: Supply and demand schedules for market M5.

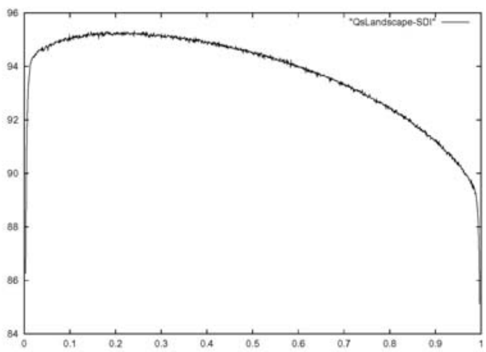

Figure 17: M4 ZI-C fitness landscape, showing allocative efficiency (vertical axis: 84.0 to 96.0) obtained for Q_s=0.0 to Q_s=1.0 (horizontal axis). Search by a GA reveals the peak fitness to be around Q_s=0.26.

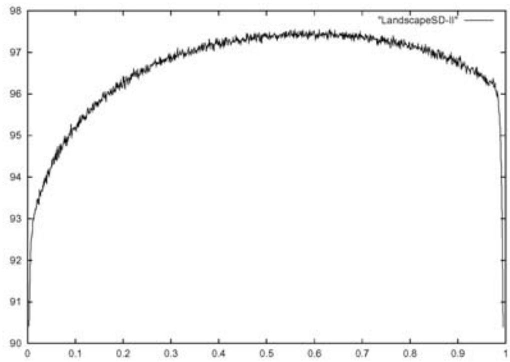

Figure 18: M5 ZI-C fitness landscape, showing allocative efficiency (vertical axis: 90.0 to 98.0) obtained for Q_s=0.0 to Q_s=1.0 (horizontal axis). Search by a GA reveals the peak fitness to be around Q_s=0.74.

V. DISCUSSION AND CONCLUSION

Section III showed the newly-calculated fitness landscapes for ZIP traders operating in marketplaces with supply and demand schedules M1, M2, M3, M1M2 and M2M1 as explored by the GA in previous papers [7,8]. All of the ZIP fitness landscapes are asymmetric (in that they do not show a symmetry around the Q_s=0.5 line, which might be expected). And several show peak performance (i.e. minimum deviation of transaction prices from the equilibrium price) at non-standard "hybrid" values of Q_s such as 0.25 or 0.16, as discovered by the GA. Although computationally expensive to calculate, these fitness landscapes reinforce the claim that in [7,8] the GA was identifying the global optima.

The primary contribution of this paper came in Section IV with the demonstration (in Figures 12, 13, 14, 17, and 18) that marketplaces populated by traders other than ZIPs surprisingly also exhibit maximal performance at non-standard "hybrid" values of Q_s. In all of the ZI-C fitness landscapes there is also a clear asymmetry, and peak performance occurs at non-standard values of Q_s. The differences in style of asymmetry and optimal value of Q_s between Figures 12 to 16 can only be attributable to the differences in the underlying market supply and demand schedules. See [10] for further exploration of this issue.

The simplicity of the ZI-C algorithm is appealing, as it offers the prospect of greater analytic tractability in attempting to identify and understand what factors in the supply and demand schedules give rise to the observed asymmetries and optima in the performance landscapes. Further research will be directed at understanding these interactions, as their nature is currently unclear. Nevertheless, the fact that the fitness landscapes for both ZIP and ZI-C traders exhibit these asymmetries and optima at non-standard values of Q_s adds strong weight to the claim that any online electronic marketplace populated by software-agent traders may give better or optimal performance if a non-standard value of Q_s is used.

While the GA has been demonstrated to be one way of identifying appropriate non-standard values of Q_s, it is essentially an off-line "batch-mode" process, poorly suited to identifying appropriate values for Q_s "online" as the market supply and demand curves alter in real-time. The search for methods that incrementally alter Q_s to appropriate values in real time while the market is "live" remains a topic for further research.

Acknowledgements

Partial financial support for Vibhu Walia came from European Commission grant ASI/B7-301/97/0126-08; further support came from the HP Labs Biologically-Inspired Complex Adaptive Systems (BICAS) research group: http://www.hpl.hp.com/research/bicas.

References

[1] D. Cliff, "Minimal-intelligence agents for bargaining behaviours in market environments". Technical Report HPL-97-91, Hewlett-Packard Labs. (http://www.hpl.hp.com/techreports/97/HPL-97-91.html)

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

[3] D. Gode & S. Sunder, "Allocative efficiency of markets with zero-intelligence traders". *Journal of Political Economy* **101**:119-137, 1993.

[4] R. Das, J. Hanson, J. Kephart, & G. Tesauro ,"Agent-human interactions in the continuous double auction". *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-01)*, 2001. (http://www.research.ibm.com/infoecon/researchpapers.html)

[5] D. Cliff, "Genetic optimization of adaptive trading agents for double-auction markets". *Proceedings of Computational Intelligence in Financial Engineering (CIFEr98) 1998.* IEEE/IAFE/Informs (preprint proceedings), pp.252-258, 1998.

[6] D. Cliff, "Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double-auction markets". Technical Report HPL-2001-99, Hewlett-Packard Laboratories, 2001. (http://www.hpl.hp.com/techreports/2001/HPL-2001-99.html)

[7] D. Cliff, "Evolution of market mechanism through a continuous space of auction-types". Presented at *Computational Intelligence in Financial Engineering (CIFEr02)* session at *WCCI2002*, Hawaii, May 2002. Technical Report HPL-2001-326: Hewlett-Packard Laboratories, 2001. (http://www.hpl.hp.com/techreports/2001/HPL-2001-326.html)

[8] D. Cliff, "Evolution of market mechanism through a continuous space of auction-types II: Two-sided auction mechanisms evolve in response to market shocks". Presented at *Agents for Business Automation*; Las Vegas, June 2002. Technical Report HPL-2002-128: Hewlett-Packard Labs, 2002. (http://www.hpl.hp.com/techreports/2002/HPL-2002-128.html)

[9] S. Phelps, S. Parsons, P. McBurney, & E. Sklar, "Co-Evolution of auction mechanisms and trading strategies: towards a novel approach to micro-economic design". Presented at *ECOMAS 2002*, New York, July 2002. (http://www.csc.liv.ac.uk/research/techreports/tr2002/tr02004abs.html)

[10] V. Walia, *MSc Thesis*. University of Birmingham School of Computer Science, forthcoming September 2002.