

Visualizing Search-Spaces for Evolved Hybrid Auction Mechanisms

Dave Cliff

Hewlett-Packard Laboratories, Filton Road, Bristol BS34 8QZ, England, U.K.
dave.cliff@hp.com

Abstract

A sequence of previous papers has demonstrated that a genetic algorithm (GA) can 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. Qualitatively similar results (i.e., non-standard hybrid mechanism designs being evolved) have been demonstrated for Cliff's ZIP trader algorithm and also for Gode & Sunder's ZI-C traders, provoking the possibility that such hybrid markets may be optimal for *any* marketplace populated entirely by artificial trader-agents. The financial implications of this work could potentially be measured in billions of dollars. In an attempt to elucidate *why* these evolved hybrid markets outperform traditional human-designed mechanisms, this paper presents results from thousands of repetitions of the GA experiments. These data allow 2-d projections of the 10-d real-space fitness landscape to be made, which *inter alia* illustrate a surprisingly high sensitivity in the relationship between the fitness evaluation function and the resulting landscape.

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 [3], 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. Originally, the values of these parameters were set by hand, using "educated guesses" [1]. However, at CIFEr'98, the first results were presented from using a standard genetic algorithm (GA) to automatically optimize

these parameter values [5,6], thereby eliminating the need for skilled human input in deciding the values of the parameters.

In all previous work using artificial traders, ZIP or otherwise, the market mechanism (i.e., the type of auction that 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); 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).

Cliff [7] presented the first results from experiments where a GA optimized 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 [7] was the first demonstration that radically new market mechanisms for artificial traders may be designed by automatic means, thereby establishing the new field of *automated market-mechanism design*. Independently,

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. This question was answered by Walia [11] who demonstrated qualitatively similar results using Gode & Sunder's ZI-C traders. ZI-C traders are essentially parameter-free stochastic agents, and their computational simplicity made it possible to explicitly plot the "fitness landscapes" showing the economic performance of ZI-C marketplaces over the entire space of possible auction mechanisms being considered. Doing the same for ZIP traders is a much more computationally daunting task, as the trader algorithm is more complex (involving a minimum of eight free real-valued parameters) and also has a richer statistical structure (i.e., it is significantly multimodal). In [10], coarse projections of the ZIP fitness landscapes were produced for comparison with the ZI-C landscapes. Here, we present the results of 4000 experiments performed to better illustrate the ZIP fitness landscapes.

The 4000 experiments reported here would represent approximately 8000 hours of continuous computation had a single-CPU computer been used (assuming the use of a 1.8Ghz P4 Hewlett-Packard ePC with 512Mb RAM running Windows2000); that is, roughly 11 months of continuous computation. Luckily, a (shared) compute-cluster built from 50 such ePCs was available for this research, reducing the minimum wait to 160(=8000/50) hours, i.e. a little under one week. Yet by the current publicly-stated aims of various computer companies, a 50-node cluster is very small fry. Super-clusters of "grid" computers or "utility data centers" are currently planned (or under construction) involving thousands or tens of thousands of such nodes. This paper concludes with discussion of the implications for artificial evolution research of such a several-orders-of-magnitude increase in available compute resources.

Before that, we give a background overview of our previously-published experiment methods and results [5,6,7,8,10]. Then we then show new visualizations that illustrate how the fitness landscape alters as a key parameter of the evaluation function is gradually changed.

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 idealized "target" price $\tau_i(t)$. When calculating $\tau_i(t)$, traders introduce a small random absolute perturbation generated from $U[0,c_a]$,¹ and also a small random relative perturbation coefficient generated from $U[1-c_r,1]$ (when a trader is reducing its $p_i(t)$) or $U[1,1+c_r]$ (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, i.e.: $\mu_i(0)=U(\mu_{min}, \mu_{min}+\mu_\Delta)$; $\beta_i=U(\beta_{min}, \beta_{min}+\beta_\Delta)$; and $\gamma_i=U(\gamma_{min}, \gamma_{min}+\gamma_\Delta)$; $\forall i$.

Hence, to initialize an entire ZIP-trader market it is necessary to specify values for the six market-initialization parameters μ_{min} , μ_Δ , β_{min} , β_Δ , γ_{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 initialization 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 further below. Before that, we discuss the issue of simulating the passage of time.

¹ Note that, throughout this paper, $v=U[x,y]$ is used to denote a random real value v generated from a uniform distribution over the range $[x,y]$.

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 [1,5,6,7,8,10]. 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 [1,5,6,7,8,10], 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 $p_j(t) \leq q(t)$ 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 equi-probably 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 either a maximum number of transactions have occurred, or until either no seller or no buyer is able to quote, or until a maximum number of time-slices have passed since the last accepted quote (i.e., a until a protracted sequence of successive ignored quotes occurs). This protracted-sequence-of-ignored-quotes termination criterion is of relevance to the discussion in later in this paper, so we should note here that in the implementation used for this work (as in the version published in [1]) the number of ignored quotes necessary to terminate trading is set by the

system constant MAX_FAILS, which has been set to 100 in the work reported here, as in all the ZIP work reported previously in [1,5,6,7,8,10].

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 IID 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 , & 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 is appealing to use an automatic exploration process, such as a GA, to identify useful Q_s values.

Thus, a 9th dimension was added to the search space, and the genotype is now the eight real values for ZIP-trader initialization, plus a real value for Q_s : now the GA is searching \mathbf{R}^9 for vectors giving the best market dynamics.

C. The Genetic Algorithm

A simple genetic algorithm was used. As with each experiment reported in [5,6,7,8,10] 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 tournament selection. Elitism (where an unaltered 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 $U[0,1]$ 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 $U[-m(g),+m(g)]$ 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 [7,8]), $N_g=1000$, $n_g=500$, $m_s=0.05$, 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 to be evolved. All genome values of μ_A , β_A , and γ_A were also clipped to satisfy $(\mu_{min} + \mu_A) \leq 1.0$, $(\beta_{min} + \beta_A) \leq 1.0$, & $(\gamma_{min} + \gamma_A) \leq 1.0$.

The fitness of genotypes was evaluated using the same methods as described in [5,6,7,8,10]. One *trial* of a particular genome was performed by initializing a ZIP-trader market from the genome, and then allowing the 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 was ended in the manner described above.

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 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 is to *minimize* the fitness value.

D. Previous Results

In [7], three differing market supply and demand schedules were used, and referred to as markets M1, M2, and M3. Only M1 and M3 will be discussed in this paper. Figure 1 shows M1: it specifies a supply and demand schedule for a marketplace with 11 buyers and 11 sellers, each allowed to buy/sell one unit of commodity, and is identical to one of the schedules used by Smith [2]; the slopes of the supply and demand curves are symmetric (i.e. gradients are equal in magnitude but opposite in sign). In M1 the difference between the limit prices of successive traders, a value we will denote by $\Delta\lambda$, is 0.25. The only difference between M1 and M3 is that, in M3, $\Delta\lambda=0.30$.




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

Figure 2 shows results from 50 repetitions of an experiment where the GA explores the \mathbf{R}^9 subspace, attempting to optimize 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 3 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 ($n=48$) while a small minority

($n=2$) settle on a less good mode of around 5.1. Clearly then, the evolved-mechanism results are better than the fixed-mechanism CDA results. When the GA is allowed to find its own value of Q_s , rather than have the CDA value of $Q_s=0.5$ imposed on it, it finds fitter solutions; solutions with less deviation of transaction prices from the equilibrium price. As it happens, $Q_s \approx 0.0$ was the value evolved in the best elite mode for the evolving-mechanism M1 experiments; but, surprisingly and significantly, for M3 the best Q_s was neither 0.0, nor 0.5, nor 1.0 (i.e. none of the Q_s values corresponding to traditional human-designed auction mechanisms); rather, the best Q_s for M3 was found to be around 0.16 [7].

To better understand the experimental results in [5,6,7,8], coarse low-resolution visualizations of the underlying fitness landscapes were computed and presented in [10]. Because 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 visualize the “fitness landscape” as a surface over the 9-d axes of the genotype space. Visualizing such a 10-d object in the two or three dimensions that we humans are familiar with communicating in is manifestly problematic; yet appropriate visualizations can be highly valuable in demonstrating that the results of the GA’s evolutionary search are indeed a plausible global optimum.

To understand the rationale for the visualization, consider Figure 3. 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 3 we would see two distinct peaks in the histogram: a big one around 4.5 and a smaller one around 5.2.

To visualize the entire fitness landscape for ZIP traders in M1, more fixed-mechanism experiments were then run, but for each set of 50 repetitions the value of Q_s was held fixed at some value while all the other 8 ZIP parameters on the genome are optimised by the GA. These data allowed 3-d projections of the 10-d fitness landscape to be plotted: in one such projection shown in [10], one horizontal dimension is the fixed value of Q_s ; another is the elite-fitness mode-value; and the vertical axis is the frequency with which the different mode values are reached for each of the

Q_s values. A perspective projection of such a 3-d histogram, calculated for ZIP traders in M1, is shown in [10]. An alternative view of the same data is also presented in [10], i.e. as a contour plot with a logarithmic compression function applied to the frequency values.




Figure 2: Elite fitness values from 50 repetitions of a 500-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 3: 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.

All of the contour plots of the fitness landscapes in [10] show good agreement with the previous experimental results, in the sense that the minima on the contour plots are in good agreement with the values discovered by the evolving-mechanism GA experiments reported on in [7,8] and summarized above: market M1 has a minimum at $Q_s=0.0$; while M3 has a minimum between $Q_s=0.125$ and $Q_s=0.25$. It is also notable that all the contour plots in [10] show a degree of multi-modality for some values of Q_s .

Thus, the evolving-mechanism results from [7,8] were supported by the brute-force exploration of the fixed-mechanism fitness landscapes for each market schedule in [10]: in each case, the evolved value of Q_s was very close to the globally optimal 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.

However, the contour plots in [10] revealed a surprisingly large difference between the landscape for market M1 and the landscape for market M3. This difference is intriguing: the two contour plots show some major differences, yet the *only* difference in the supply and demand schedules is that in M1 $\Delta\lambda=0.25$ while in M3 $\Delta\lambda=0.30$. Specifically, in both M1 and M3 the equilibrium price is 2.00 and the equilibrium quantity is 6; both schedules have 11 buyers and 11 sellers; and in both schedules the $\Delta\lambda$ value is constant and identical for all buyers and for all sellers. That such a minor difference in the market schedules can have a relatively major effect on the resultant fitness landscapes is an issue deserving a detailed exploration, and that is the topic of the rest of this paper.

Having established the background to our current work, we now proceed with introducing our new results. In the next section we show finer-grained projections of the fitness landscapes that result from using schedules M1 and M3, and also from intermediate schedules where the step-size $\Delta\lambda$ has values between 0.25 and 0.30.

High-Resolution ZIP Fitness Landscapes

A. Methods

The contour plots presented in [10] are one approach to visualizing densities of the elite-fitness modes, but the fact that the data was calculated at fixed discrete Q_s values (of 0.000, 0.125, 0.250, ..., 1.00) introduced some unappealing artifacts. For instance, smooth continuous features in the landscape can appear jagged or even discontinuous as a result of the discretization of the range of Q_s values. In principle, a finer-grained regular sampling of Q_s values (e.g. at $Q_s=0.000, 0.005, 0.010$, etc) would have provided data to reduce these artifacts, but another approach, explored here, is to use random jitter in the choice of Q_s values.

Another problem with the contour plots of [10] is that they show histograms for the elite-fitness data at generation 500, but give no indication of the variability in that data (e.g., how different would the contour plot have looked at generation 499, or 475?). It would be more rigorous and more informative to show measures of central tendency (e.g. a mean) and of variability (e.g. a standard deviation).

To avoid these visualization artifacts, and to show the data in a more rigorous fashion, a new set of visualization experiments was conducted. In each experiment, a fixed value for Q_s was chosen at random from a uniform distribution, and that value remained constant throughout

the experiment. All other experiment details were the same as those described above, except that the final measure of an experiment's outcome is now the mean and standard deviation of the elite fitness over the final 50 generations of the experiment (i.e., generations 450 to 500). The result of any one such experiment can be plotted on a 2-d graph showing Q_s on the horizontal axis, with the vertical axis showing mean elite fitness (and error bars indicating the standard deviation on that mean). Performing a large number of repetitions of this experiment (e.g. $n=500$) with a different randomly-generated value of Q_s on each repetition gives a scatter-plot which as a visualization technique is broadly similar to the contour plots shown previously in [10] but which avoids their problems. In the remainder of this paper, f_μ will denote mean elite fitness.

B. Results

Q_s/f_μ scatter-plots were generated for $\Delta\lambda=0.25, 0.26, 0.27, 0.28, 0.29$, and 0.30 . They are illustrated in Figs 4 to 9 respectively. The data for M1 ($\Delta\lambda=0.25$) in Fig. 4 shows three clear modes. A dominant mode runs from $f_\mu\sim-4.1$ at $Q_s=0.0$ through to $f_\mu\sim-4.3$ at $Q_s=0.5$. For later discussion, we'll call this Mode A. Above and roughly parallel to Mode A there lies a second weaker mode running from $f_\mu\sim-4.75$ at $Q_s=0.0$ to $f_\mu\sim-5.00$ at $Q_s=0.5$, which we'll call Mode B. Finally, a third mode sweeps up and to the left from $f_\mu\sim-3.5$ at $Q_s=0.00$ up to $f_\mu\sim-5.5$ at $Q_s=0.25$, beyond which it rapidly fades. We'll call that Mode C.

Examining Figures 5 to 9, we see that each change of 0.01 in $\lambda\Delta$ has a noticeable systematic change on Modes A, B, and C. At a macro-level, we see that Modes B and C act as less powerful attractors as $\Delta\lambda$ increases, so that by $\lambda\Delta=0.30$ (i.e., M3) Mode A appears to be effectively all that remains. At the micro-level, in the change from $\lambda\Delta=0.25$ to $\lambda\Delta=0.26$, Modes A and B have both moved up slightly, while Mode C has moved upwards more rapidly. Furthermore, for $Q_s=0.25$ to $Q_s=0.50$, a new mode (which we'll call mode D) is appearing slightly above mode A. The distinction between Modes A and D is seen more clearly when the error bars are removed from the scatter plots (as presented in [12]) which helps to show that, as $\lambda\Delta$ increases, Mode D becomes a stronger attractor; such that by $\lambda\Delta=0.28$ the dominant mode is actually D rather than A.

It so happens that the distinct f_μ modes are clearly (and conveniently) strongly correlated with the μ_{min} (base profit margin at time zero) locus on the genotype, allowing distinct clusters to be identified by inspection. A series of scatter-plot figures presented in [12] show the μ_{min}/f_μ data for $\lambda\Delta=0.25$ to 0.30 respectively. These μ_{min}/f_μ plots show that genomes in Mode C all have $\mu_{min}\sim-1.0$ and that as $\lambda\Delta$ increases the influence of Mode C weakens, failing to

attract any of the elite genotypes once $\lambda\Delta > 0.28$. Mode B can also be seen as a distinct cluster with μ_{min} values spread from ~ 0.0 to ~ 0.25 ; the f_{μ} values for this cluster rises and fades as $\lambda\Delta$ increases.




Figure 4: $\lambda\Delta=0.25$ Q_s/f_{μ} scatter-plot. Vertical axis is fitness, horizontal axis is fixed Q_s value. Each data-point marks the mean elite fitness f_{μ} from the final 50 generations of a 500-generation experiment, with the vertical error bars indicating plus and minus one standard distribution. This plot shows data from 500 repetitions of the experiment with the fixed Q_s value for each experiment being generated from $Q_s=U[0.0,0.5]$.




Figure 5: $\lambda\Delta=0.26$ Q_s/f_{μ} scatter-plot. Format as for Figure 4.




Figure 6: $\lambda\Delta=0.27$ Q_s/f_{μ} scatter-plot. Format as for Figure 4.




Figure 7: $\lambda\Delta=0.28$ Q_s/f_{μ} scatter-plot. Format as for Figure 4.




Figure 8: $\lambda\Delta=0.29$ Q_s/f_{μ} scatter-plot. Format as for Figure 4.




Figure 9: $\lambda\Delta=0.30$ Q_s/f_{μ} scatter-plot. Format as for Figure 4.

The interplay between Modes A and D as $\lambda\Delta$ increases is made more clear when the same co-ordinate data is normalized by the underlying value of $\lambda\Delta$: that is, when the mean elite fitness value and the mean elite μ_{min} value for each data-point are divided by the $\lambda\Delta$ for the experiment. This normalization accounts for the fact that μ_{min} is interpreted as a *percentage* margin. The normalized data (plotted in a sequence of graphs in [12]) shows that Mode A is characterized by $\mu_{min}/\lambda\Delta > 1.0$ & $16 < f_{\mu}/\lambda\Delta < 18$; while Mode D is characterized by $\mu_{min}/\lambda\Delta < 1.0$ & $17 < f_{\mu}/\lambda\Delta < 19$. And so it is clear that the dominant mode for $\lambda\Delta=0.30$ is one that acts as only a very weak attractor for $\lambda\Delta=0.25$

(with the hindsight gained by this analysis, Mode D can now be seen as the scattering of nine data-points above Mode A between $Q_s=0.375$ and $Q_s=0.500$ in Figure 4).

Having clarified the nature of the four modes seen in the fitness data of Figures 4 to 9, we can turn our attention to the nature of the best mode in each experiment. For $\lambda\Delta$ in the range $[0.26,0.30]$ we see in Figures 5 to 9 that, for both Modes A and D, there is a clear shallow “U” curve with its minimum f_μ (i.e., most desired market dynamics) at Q_s values around 0.125. This is true also of the Mode A curve for $\lambda\Delta=0.25$, but in that instance we see also that Mode C actually gives lower fitness (i.e., better market dynamics). Given the upward path of Mode C identified as $\lambda\Delta$ increases, we can expect to see Mode C acting as a much stronger attractor for $\lambda\Delta < 0.25$. Given the large differences in genome μ_{min} and optimal Q_s values between Mode C and Modes A/D, it would appear that $\lambda\Delta=0.25$ marks the point of a “phase transition” between two radically different modes of elite genotypes.

To confirm the existence of this transition, two more sets of 500 experiments were conducted, for $\lambda\Delta=0.24$ and $\lambda\Delta=0.23$. The respective Q_s/f_μ and μ_{min}/f_μ scatter plots are presented in [12]. The increased dominance of Mode C as $\lambda\Delta$ reduces is clear from these extra data. The sudden transition of the elite genotype mode from Mode A/D to Mode C is masked by a smooth transition in f_μ values, as illustrated in Figures 10 and 11, which respectively show the f_μ values and the elite-mode μ_{min} values, each as a function of $\lambda\Delta$, for the best-scoring 1% ($n=5$) of the 500 repetitions for each value of $\lambda\Delta$. While Figure 10 shows a progressive increase (i.e., a worsening) in the fitness of the best 1% as $\lambda\Delta$ increases, Figure 11 illustrates the “phase transition” step-change in μ_{min} values at $\lambda\Delta=0.25$.

C. Discussion

The results presented in this section, from 4000 repetitions of the GA experiments, are more illuminating than the previous contour-plot visualizations of these fitness landscapes first published in [10]. Four distinct elite-genotype modes have been identified, and the effects of changes in $\lambda\Delta$ on those modes have been revealed. In several places, a “snapshot” approach, i.e. studying f_μ distributions for a single fixed value of Q_s , (as used in the contour plots of [10]) would give apparently bi-modal results where in fact the data is tri-modal but two modes overlap in the Q_s/f_μ scatter-plot: this is seen most clearly in Figure 4 where Modes A and C overlap for much of the range $Q_s=0.0625$ to $Q_s=0.3125$.

While it is fortunate that the four modes reveal themselves so clearly as distinct clusters in the μ_{min} data, if this had not

been the case then any appropriate clustering algorithm could be used to automatically identify such clusters. Similarly, it is fortunate that the single Q_s dimension yielded such revealing scatter-plots. But in the absence of such good fortune, there are many techniques for dimensionality reduction in multivariate data that could be employed to identify the best dimensions to use as basis vectors for projecting down from the high-dimensional source data to the two or three dimensions appropriate for display on screen or on paper.




Figure 10: $\lambda\Delta/f_\mu$ scatter-plot: horizontal axis is $\lambda\Delta$, vertical axis is f_μ ; data-points are from the best-scoring 1% ($n=5$) of the 500 repetitions performed for each value of $\lambda\Delta$.




Figure 11: $\lambda\Delta/\mu_{min}$ scatter-plot: horizontal axis is $\lambda\Delta$, vertical axis is μ_{min} ; data-points are from the best-scoring 1% ($n=5$) of the 500 repetitions performed for each value of $\lambda\Delta$.

The predominant change resulting from successive reductions in $\lambda\Delta$ from 0.30 to 0.23 is the increasing influence of Mode C as an attractor for elite genotypes. The visualization of the search space has demonstrated that $\lambda\Delta=0.25$ marks the point of a phase transition where the fittest genotypes found in an evolving- Q_s GA experiment cease to come from Modes A/D and start to come from Mode C. At face value, this would indicate that Mode C genotypes are the best to use for markets such as M1 where $\lambda\Delta \leq 0.25$. However, this is not necessarily the case. As is discussed in [12], Mode C genotypes score a high fitness by exploiting an aspect of the ZIP implementation not

considered in the fitness evaluation function. Recall that earlier in this paper it was stated that in the original ZIP implementation published in [1] (and used in all successive ZIP work [5,6,7,8,10]) there is a system parameter `MAX_FAILS`, set to 100, that determines the maximum number of successive ignored quotes allowed to pass before a trading period is ended. This parameter prevents ZIP marketplaces from continuing indefinitely with traders constantly quoting prices ignored by the contraside (and making small adjustments to their margins as a consequence) but never actually trading.

As is discussed in more depth in [12], the Mode C genotypes exploit the fact that, with `MAX_FAILS=100`, very long sequences of quotes can occur without any transactions taking place: for Mode C the bids announced in the market during the protracted periods of non-trading allow the traders to adjust their margins such that, *when transactions do start to occur*, equilibration is significantly quicker. And so we see that the Mode C data is a reflection of the old adage “you get what you measure”. The fitness evaluation function used here (and in all our previous work) is dependent only on *transaction* prices, and totally ignores the *number of ignored quotes* in the market leading up to each transaction. This is exploited by the best Mode C genotypes, which specify ZIP initialization parameters that collectively ensure a high likelihood of *both* many failed bids preceding a transaction *and* a bid actually being accepted before the `MAX_FAILS` limit is reached, where the seller accepting the bid is likely to be the most intra-marginal seller remaining in the market.

One obvious way around this type of solution would be to employ a fitness evaluation function that rewards higher ratios of accepted-to-ignored quotes as well as rewarding equilibration speed. As these are two distinct objectives, a Pareto-optimizing approach would be appropriate.

Although the results presented in this paper (and in [12]) have illuminated the nature of the surprising difference in the fitness-landscape contour-plots for M1 and M3 shown in [10], and although the discussion in [12] identifies how the solution encoded on Mode C genomes operates, several significant open questions remain to be answered: these are enumerated in [12].

All of those questions appear to be fertile ground for future research, and for the time being the most productive way of tackling them appears to be by empirical exploration methods such as those presented here. The results shown in this section make clear that small variations in one free parameter (i.e., the $\lambda\Delta$ value affecting the simple stationary symmetric supply and demand schedules in the ZIP-trader marketplaces) can have effects that are not intuitively predictable in advance. This lack of an intuitive

understanding could in principle be addressed by appropriate mathematical modeling, but the compounded non-linearities, stochasticity, and adaptation (giving non-stationary probability distributions) in the ZIP-trader system all combine to make a full analytic understanding (e.g., via game-theoretic analysis or probabilistic modeling) seem a very distant goal. The chicken-and-egg circularity in the relationship between theory and data is well known; for the time being, we are concentrating on building a body of parametric-variation data that can then be used to suggest and constrain hypotheses/theories that guide the direction of subsequent empirical studies.

Clearly, there is no shortage of topics for further research; and nor is there any shortage of experiments with which to keep our compute-clusters busy.

Conclusion

This paper commenced with a review of previous experimental work up to and including the recent “low-resolution” contour plots that give meaningful 2-d projections of the 10-d fitness landscapes underlying this work. The major new results of the previous section showed the outcome of 4000 repetitions of our GA experiments (taking ~8000 hours of computation).

This higher-resolution data showed that, with only very minor variations in the $\lambda\Delta$ value, the “attractor” modes for the elite genotypes in our experiments could show major shifts and phase transitions. While the $\lambda\Delta$ value is not an *explicit* term in the fitness evaluation function, it is an *implicit* factor insofar as it determines the supply and demand curves that underlie the equilibration-based fitness evaluation function used in this work; and of course the evaluation function plays a large part in determining the fitness landscape for any GA system.

The new high-resolution data again confirms that the GA can successfully find optimum genotypes. However, the data on *sub-optimal* solutions (i.e. the nature of the fitness landscape at points *other* than optimum) has the potential to be more informative than the optimum data taken alone. Although none of the evolved solutions here improve significantly on any of our original GA results presented in [6], the new data helps us to better illuminate the nature of the fitness landscapes traversed by the GA in attempting to evolve new mechanism designs for ZIP-trader marketplaces, and also helps us to understand the nature of the optima in those landscapes. With this illumination comes the chance of better understanding *what* factors in the system affect the nature of the optimum solutions, and *how* they do so: at the moment, the causal mechanistic interactions leading to one genome being better than another for

any particular supply/demand schedule remain unclear, rendering all our results as idiopathic.

While the high-resolution visualizations presented for the first time here offer better opportunities for understanding the performance of the evolving-marketplace ZIP-trader systems, it should be noted that none of the figures in this paper are at all novel in terms of the nature of the graphics techniques employed: scatter plots are nothing new.

But what is perhaps more novel here is the approach to consuming CPU cycles: burning 8000 hours of CPU time to generate these figures would have seemed an astonishingly profligate use of compute resources only a few years ago. And indeed, for a researcher working with a single-CPU machine, waiting eleven or so months for a set of possibly-illuminating visualization results would perhaps not be the best use of computer time.

Yet, with many hardware vendors now pushing to build “utility data centers” and/or “grid” compute facilities with thousands or tens of thousands of connected processors available offering CPU cycles as a utility-style resource, the shared 50-server compute farm used for this work starts to appear decidedly lightweight, despite the fact that it can be spoken of as being roughly equivalent to a single-CPU 90Ghz Pentium4 machine with 25Gb of RAM. Within the next few years, it seems perfectly plausible that research such as the explorations reported here could be conducted by buying CPU-cycles from a utility data center (UDC) in such a way that the time taken to generate the visualization data is measured in minutes or even seconds, rather than days or weeks. In such a scenario, the transition in working styles will be similar to the move from *batch-mode* to *interactive* computing that was brought about by the development of mini/microcomputers in the 1960s and 1970s. The need would then be not so much for newer and more efficient or clever ways of getting a GA to find a solution, but rather for faster and more informative ways to visualize the high-dimensional data streams pouring out of the UDC; and for more intuitive ways to steer through these data spaces to find areas of significance or interest. In fact, research interest in “blind” search techniques such as GAs and other forms of artificial evolutionary computation could possibly wane, to be replaced by a new style of working involving applications of brute-force enumerative search to spaces of possible solutions so large as to be considered practically infinite, where the search is guided interactively by skilled human operators interacting with sophisticated visualization workstations (where “visualization” could include presentation via non-visual sensory modes such as audio, force-feedback on controllers, etc). Right now this seems a realistic and exciting possibility, but only time will tell.

Acknowledgements

This paper is abridged from [12]. Thanks to Andrew Byde for valuable discussions, and to the reviewers for helpful comments. Computing resources came from the HP Labs 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, HP Labs, 1997. (<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”. *Proc. 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”. *Proc. Computational Intelligence in Financial Engineering (CIFEr98) 1998*. IEEE/IAFE/Informs, 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 ‘02. 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*, NY, July 2002. (<http://www.csc.liv.ac.uk/research/techreports/tr2002/tr02004abs.html>)
- [10] D. Cliff, V. Walia, & A. Byde, “Evolved Hybrid Auction Mechanisms in Non-ZIP Trader Marketplaces” Technical Report HPL-2002-247, Hewlett-Packard Laboratories, 2002. (<http://www.hpl.hp.com/techreports/2002/HPL-2002-247.html>)
- [11] V. Walia, *Evolving Market Design*. MSc Thesis, University of Birmingham School of Computer Science, September 2002.
- [12] D. Cliff, “Visualizing Search-Spaces for Evolved Hybrid Auction Mechanisms” Technical Report HPL-2002-291, HP Labs, 2002. (<http://www.hpl.hp.com/techreports/2002/HPL-2002-291.html>)