Aliasing in XCS and the Consecutive State Problem : 1 - Effects

Alwyn Barry

Faculty of Computer Studies and Mathematics,
University of the West of England,
Coldharbour Lane, Bristol, BS16 1QY, UK

Email: Alwyn.Barry@uwe.ac.uk
Phone: (++44) 117 965 6261 ext. 3777

Abstract

Whilst XCS (Wilson, 1998) has been shown to be more robust and reliable than previous LCS implementations (Kovacs, 1996, 1997), Lanzi (1997) identified a potential problem in the application of XCS to certain simple multi-step non Markovian environments. The 'Aliasing Problem' occurs when the environment provides the same message for two states in environmental positions that generate different constant payoffs. This prevents classifiers forming a correct payoff prediction for that message. This paper introduces a sub-class of the aliasing problem termed the 'Consecutive State Problem' and uses the subclass to identify the effects of consecutive state aliasing on the learning of the State × Action × Payoff mapping within XCS. It is shown that aliasing states can prevent the formation of classifiers covering preceding states due to the trade-off of accuracy for match set occupancy made by the classifiers covering the aliasing states. This can be prevented by identifying a condition encoding which makes such match set 'piracy' improbable. However, under conditions of intense competition for population space where the classifier covering the aliased states cannot gain additional match set occupancy these classifiers will not be maintained within the population. Barry (1999) uses these findings to identify a solution to the Consecutive State Problem which is less heavyweight than the more general solution proposed by Lanzi (1997, 1998).

1 INTRODUCTION

XCS (Wilson, 1995, 1998) is a Learning Classifier System (Holland, 1986) with ancestry in the Animat and ZCS LCS implementations (Wilson, 1983, 1994). It maintains the basic condition-action structure and ternary/binary encoding of the 'traditional' LCS with novel mechanisms for recording the 'strength' of the classifier which separate the measure of performance utility within a given situation (payoff 'prediction') from that of reproductive utility ('fitness', based on predictive accuracy). By adding GA niching mechanisms derived from work by Booker (1982) XCS achieves the ability to discover and maintain a full State × Action × Payoff mapping for the test environment, with optimal levels of generalization (Kovacs, 1996). XCS thus represents a quantum leap forward in the reliability and utility of Learning Classifier Systems. The interested reader is directed to Kovacs (1996) which provides a detailed explanation of the construction and operation of XCS.

Early results for the application of XCS to simple test environments were presented by Wilson (1995, 1998) and Kovacs (1996, 1997). These were predominately focused on single step environment, although Wilson presented results for XCS within the Woods 2 environment (Wilson, 1995). Lanzi (1997, 1998) sought to apply XCS to more complex Markovian and non-Markovian Woods-based test environments in order to investigate multiple step environments further. As part of this work Lanzi (1998) identified a significant problem for XCS learning in multi-step environments - the aliasing of states. Within certain Woods environments it is possible to derive an input vector which is repeated elsewhere in the environment with a different payoff value. For example, in the simple Woods environment shown in Figure 1, under the payoff, parameterization, and encoding used in Wilson (1995), the two blank positions in the center of the environment will each generate the same input message 0100000100100100000010010, but the expected payoff to the right central position will be 504.1 and to the left central position will be 357.911. Since a single classifier will represent these two positions, the payoffs that this classifier will receive will vary and therefore the classifier will be adjudged to be inaccurate.

```
OOOOOO        O = Rock
O....F        . = Space
OOOOOO        F = Food
```

Figure 1 - A simple aliasing environment.

Lanzi attempted to overcome this problem by introducing a state memory mechanism proposed by Wilson (1994, 1995) and originally applied within ZCS by Cliff & Ross

(1994). Lanzi (1997) demonstrated that the mechanism was able to disambiguate internal states with aliased input within the Woods101 environment. In Lanzi (1998) this mechanism was identified as imperfect because it could generate the same memory configuration for the two aliased positions. He therefore introduced two modifications which link the internal memory register setting more closely to the external actions of the Animat and the provision of payoffs (see Lanzi (1998) for further details). He demonstrated that this mechanism is sufficient to disambiguate the environment and therefore cause separate classifiers to be generated for each of the aliased inputs.

The 'Aliasing Problem' was not central to Lanzi's research program and therefore his work has not investigated further the effects of aliasing on the classifiers covering the aliasing states or learning within the XCS as a whole. Furthermore, he has not distinguished between the varieties of state aliasing which may be found within test environments, and therefore does not identify possible solutions to these sub-classes of the Aliasing Problem which may be simpler to implement. In this paper the 'Consecutive State Problem' is identified as a sub problem of the Aliasing Problem and is used to identify some effects of the aliasing states on the performance of classifiers within various Finite State World XCS test environments. Barry (1999) extends this work to demonstrate two solutions to the Consecutive State Problem which, while not addressing the whole Aliasing Problem, are simpler to implement.

2 HYPOTHESIS

In the domain of Finite State Worlds (Grefenstette, 1987; Riolo, 1987), consider the FSW (which we will denote *FSW-5*) consisting of a start state s_0, further states s labeled s_1 to s_3 and a terminal state labeled s_4. Each of states s_0 to s_3 are sources of directed edges drawn so that for all states s_i ($0 \leq i \leq 3$) a single edge emanates from s_i and terminates in s_{i+1}. Each edge e is labeled $e_i = i + 1$ ($0 \leq i \leq 3$), which is the action required to traverse that edge . Every state emits a signal d capable of unambiguous sensory detection such that for all states s_i $d = i$. The start state is s_0 and upon reaching s_4 a reward R is given and the FSW is reset to s_0. Four classifiers are required to traverse this FSW : $0\rightarrow1$, $1\rightarrow2$, $2\rightarrow3$, $3\rightarrow4$. On reaching s_4 classifier 4 receives a reward R. Since XCS uses the Widrow-Hoff update mechanism, over successive trials classifier 4 will converge to the prediction $p = R$, and classifiers 1 to 3 will converge to the predictions $\gamma^{3-i}.R$ where γ is the discount factor applied to reduce the prediction value paid from classifiers in the current Action Set [A] to the previous Action Set [A-1].

Now consider a modification to this FSW (denoted *FSW-5A*) such that s_1 and s_2 emit the signal d=1 with the edges $s_1\rightarrow s_2$ and $s_2\rightarrow s_3$ labeled 2. Three classifiers are now required to traverse this FSW: $0\rightarrow1$, $1\rightarrow2$, $3\rightarrow4$. Again, upon reaching s_4 classifier 3 receives a reward R. Classifier 3 will converge so that $p = R$ in the same manner as Classifier 4 within *FSW-5*. Let us assume that the prediction of classifier 3 has converged to this value. For moving from s_2 to s_3, classifier 2 will be consistently given a payoff γR. However, for moving from s_1 to s_2 classifier 2 will also be given a payoff which is γP_2. If the learning rate β within the Widrow-Hoff mechanism was 1, then the prediction would oscillate well within the limits $\gamma^2 R$ and γR. For simplicity, let us assume that P_2 varies around the average payoff that would have been received at the states had they not been aliased $((\gamma R + \gamma^2 R) / 2)$, and that the learning rate β is less than unity. In this case the variance will reduce to $\pm \beta((\gamma R - \gamma^2 R) / 2)$. Unless the value of β is very small, or the aliased states are sufficiently far from the reward source for the successive application of the discount factor γ to reduce the payoff to a very small amount, the variance will remain sufficient to produce an oscillation in P_2 which is greater than ε_0, the minimum error for a classifier to be considered accurate.[1]

The 'Consecutive State Problem' is introduced to label this form of the Aliasing Problem:

Hypothesis 1
The aliasing problem is not restricted to independent states which exist at separate locations within an environment, but will also be seen whenever two or more consecutive states admit to the same sensory perception (given the limitations of the sensory system of a given Animat) and together lead to a later consistent reward.

Consider the stage in the operation of an XCS within *FSW-5A* where the payoff to the last aliased state s_3 will be constant at γR. At the start of the next trial classifier 1 moves the FSW to state s_1. At this iteration (which we will call iteration i) the prediction of the classifier covering the aliasing states is P_i. On the next iteration the prediction P_{i+1} of classifier 2 (the aliasing classifier) becomes $P_i + \beta(\gamma P_i - P_i)$, causing the classifier's prediction to reduce towards $\beta\gamma P_i$. In the following iteration P_{i+2} becomes $P_{i+1} + \beta(\gamma R - P_{i+1})$ which increases P_{i+1} towards $\beta\gamma R$. The whole update has the effect making $P_i = P_{i+2}$. The preceding classifier therefore receives a constant prediction value so long as all the aliased states are visited within each trial.

Now consider the case where each aliased state may not be visited within each trial. In *FSW-5*, the payoff delivered to classifier 1 would be $\gamma^3 R$. Within the circumstances described in *FSW-5A*, classifier 1 receives payoff from classifier 2 whose maximum prediction oscillation limits lie between $\gamma^2 R$ and γR, as discussed above. Classifier 2 will converge towards γR if in all preceding trials the trial started from s_2, but will converge to below but near $((\gamma R + \gamma^2 R) / 2)$ if in all preceding trials the trial starts from s_0 or s_1. Therefore the payoff received by classifier 1 will oscillate and the prediction of classifier 1 will oscillate. Although the learning rate β will

[1] Of course, the value P_2 will not actually vary around $((\gamma R - \gamma^2 R) / 2)$ due to the Widrow-Hoff update mechanism adopted by XCS (Wilson, 1995). The value P_2 will be below the average of the feedback to the equivalent non aliased states because the first update in each transition past s_1 and s_2 will, in effect, be averaging P_2.

reduce the degree of variance in the error in prediction calculated by the preceding classifier, for large payoffs received within aliased states late in the payoff chain and where ε_0 is small, the variance in payoff may be sufficient to cause inaccuracy in the preceding classifier.

Hypothesis 2

The classifier covering the non-aliased state immediately preceding the aliased states will be able to achieve an accurate payoff prediction in cases where each aliased state is visited in each trial, but can be considered inaccurate in cases where the aliased states are not visited within each trial.

What is the likely effect of aliasing upon the induction mechanisms within XCS? The XCS selects classifiers for reproduction using their fitness, which is based upon their relative accuracy. If we accept Hypothesis 1, the classifier which covers the Aliasing States will have a very low accuracy and therefore a fitness which is similar to other competing classifiers within each [A] it participates in. Since the classifiers covering non-aliased states (with the exception of those covering the immediately preceding states - see Hyp. 2) will eventually be classified as accurate, these will have a high fitness and be selected by the GA proportionally more often. Their numerosity will then increase, putting pressure on all inaccurate classifiers. Ultimately the combination of selection and population pressure will eradicate the classifier covering the Aliasing State. The Covering Operators will rapidly replace this classifier with another, but no replacement will be deemed more accurate.

Hypothesis 3

The aliasing of consecutive states will generate inaccuracy in any classifier that matches the sensory input and moves the Animat to the next aliasing state. The inaccurate classifier will rapidly be replaced by the action of the GA without any suitable replacement available to generate a greater degree of accuracy. This will prevent the formation of an accurate State $\times$ Action $\times$ Payoff mapping and lead to the perpetual ineffectiveness of the classifier population if no alternative set of actions is available.

3 EXPERIMENTAL INVESTIGATION

3.1 THE TEST ENVIRONMENT

Both Wilson and Lanzi have utilized the Woods environments in their work, but these environments are not easily scaled with fine control in either length or complexity. Therefore, the Woods environments are set aside in favor of a FSM-like environment similar to that proposed by Grefenstette (1987) and used extensively by Riolo (1987). A 'Finite State World' is an environment that is modeled as a finite Markov process. Similar in appearance to a State Transition Machine, each state in the environment is associated with a labeled node within a directed graph. The nodes depict the individual states, and the directed edges denote the legal transitions between states, each labeled with the action which causes an 'instantaneous' transition along the edge. Finite State Worlds can be created which are equivalent to Woods environments but FSW are more precise than Woods environments. Each state can be given a distinct label so that it is possible to ensure that aliasing problems do not occur even in long chains. Furthermore, configurations which would not be possible with a Woods environment can be created, as figure 2 demonstrates.

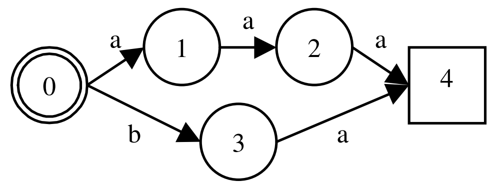

Figure 2: An Example "Woods-impossible" FSW.

3.2 HYPOTHESIS 1

To empirically verify Hypothesis 1 the *FSW-5* environment was constructed and five classifiers inserted into the initial population to cover each of the states within the environment. The XCS was allowed to run without induction algorithms and with parameterisation set to N=400, p_1=10.0, ε_1=0.01, f_1=0.01, R=1000, γ=0.71, β=0.2, ε_0=0.01, α=0.1, θ=25, X=0.8, μ=0.04, P(#)=0.33, s=20 (see Kovacs (1996) for a parameter glossary). To measure the error within the population a new measure termed 'System Relative Error' was computed. The error measure used by Wilson (1995, 1998) only captures the absolute error in the System Prediction. This work required a measure of the error in each [A] formed during each trial which accounted for the payoff discount through the payoff chain. This measure was constructed by averaging the magnitude of the error in System Prediction ((P_{i-1} - payoff) / P_{i-1} where P_{i-1} > payoff, (payoff - P_{i-1}) / payoff otherwise) for [A $_{i-1}$] during each exploitation trial and for [A] at the end of the trial, reset at the start of a new trial. Alongside this, the maximum and minimum magnitudes in the error of the System Prediction were recorded so that a measure of spread in error was also available. As figure 3 shows, the System Relative Error reduces rapidly to zero within the non-aliasing *FSW-5* in a typical run with the aforementioned parameterisation.

The environment *FSW-5A* was then constructed and the four classifiers required for this environment inserted into the population before running XCS, again without induction algorithms and with the same parameterisation. Figure 4 shows that System Relative Error fails to fall. The predictions of the classifiers within these tests were captured and plotted to identify the source of the error, and Figure 5 clearly demonstrates the oscillation within classifier 2 of the XCS population for *FSW-5A* showing this classifier to be the source of the error. These findings confirm Hypothesis 1.

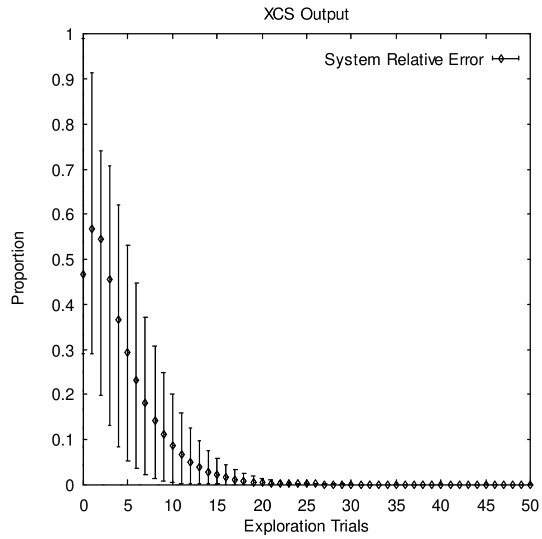

Figure 3: The decline in System Relative Error in *FSW-5*.

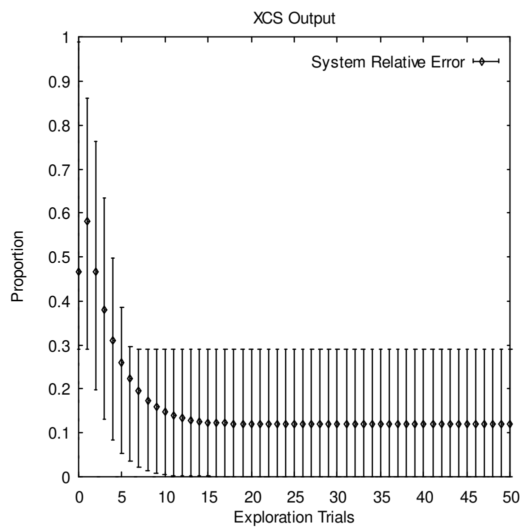

Figure 4: System Relative Error fails to fall in *FSW-5A*.

Further investigations were conducted to identify the effect that the number of aliasing states has upon the degree of oscillation in the prediction of the classifier covering the aliasing states. For these experiments the *FSW-5A* environment was extended to nine states, with the aliasing states for the 2 alias test (*FSW-9A-2*) being s_5 and s_6, expanded to the states s_3 to s_6 for the four state test (*FSW-9A-4*), and expanded further to states s_1 to s_6 for the six state test (*FSW-9A-6*). For each test the initial classifiers required by the environment were inserted into the population and XCS was run with the same parameterization as given above. Figure 6 plots the prediction of the classifier covering the aliasing states in each case, and demonstrates that an increase in the number of consecutive aliased states will increase the range of oscillation in the prediction of the classifier which covers those aliased states, as would be expected. It can also be seen that, as predicted, the stable prediction of the classifier covering the aliasing states oscillates about a lower prediction than the average discounted payoff over the aliasing states since the classifier is feeding back the discount of its moving average of the payoffs.

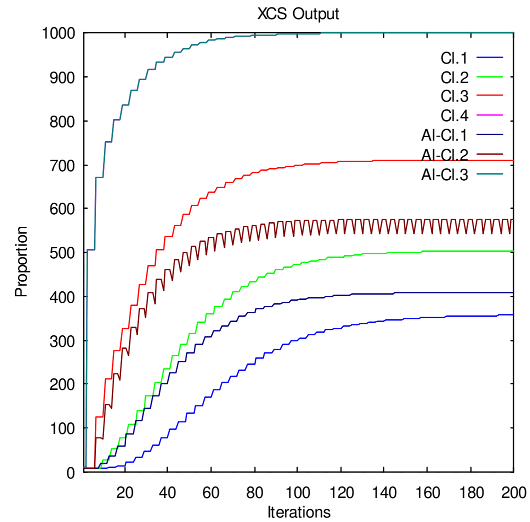

Figure 5 - Classifier predictions in *FSW-5* and *FSW-5A*.

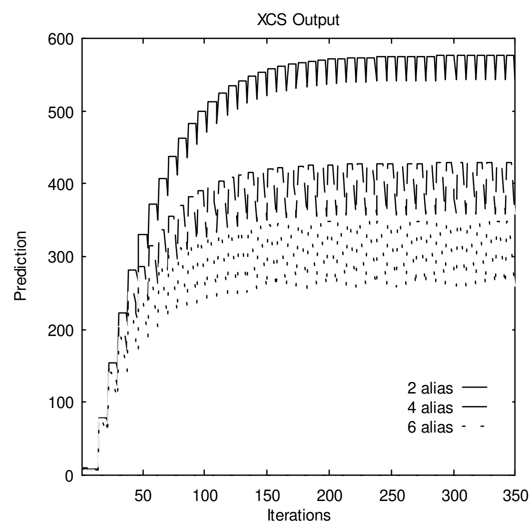

Figure 6 : Change in oscillation with more aliasing states.

3.3 HYPOTHESIS 2

The results of the experiments used to empirically prove Hypothesis 1 can be applied to address Hypothesis 2. Examining figure 5 we see that the prediction of classifier 1 has been changed to a higher prediction within *FSW-5A* than was the case for classifier 1 within *FSW-5*, as would be expected. However, the payoff given to the preceding classifier does not oscillate, indicating that the fixed point prediction of the classifier covering the aliased states remains stable at the payoff point as predicted. In order to verify the second proposition of Hypothesis 2, the experiments with 2, 4, and 6 consecutive aliased states were repeated with the same parameterization, but allowing states s_0 to s_7 to be start states with the start state chosen arbitrarily from the available start states at the beginning of each trial. Figure 7 plots the predictions from a typical run within each environment (a typical run is shown because averaging the 30 runs within each experiment hides the fluctuations), and illustrates that under these new conditions the aliasing states do affect the stability of the prediction of the preceding classifier.

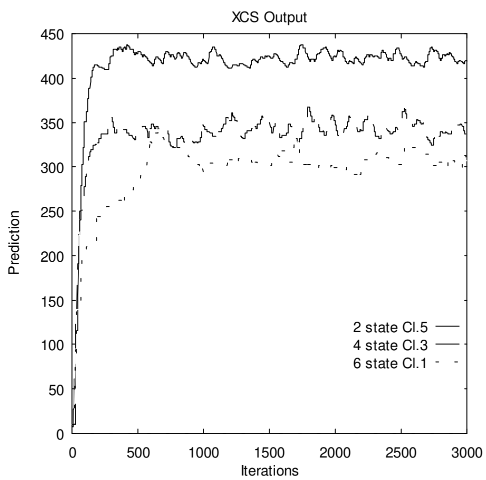

Figure 7 : The change in oscillation within the prediction of the aliasing classifier with an increased number of aliased states.

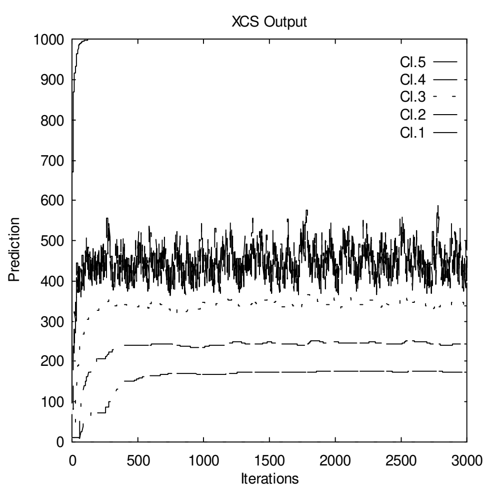

Figure 8 : All classifiers are affected by the aliasing states when the exploration rate is not uniform across all states.

Figure 8 plots the predictions of all classifiers in a typical run within *FSW-9A-4*, and demonstrates that the oscillation in prediction not only affected the immediately preceding classifier, but also influenced earlier classifiers. However, at no time were these earlier classifiers considered to be inaccurate - the oscillations had been sufficiently smoothed out by the discounting within the Widrow-Hoff mechanism to move the changes in prediction within the 1% accuracy boundary used within these problems. In all 30 runs in each environment these effects were repeated and therefore it is concluded that Hypothesis 2 is verified by these findings.

3.4 HYPOTHESIS 3

To obtain a baseline performance a two action nine state FSW termed *FSW-9(2)* was used, with the first action causing a transition to the following state and the second action causing a transition to the current state - an effective null action. All parameterization was as given previously. The GA and covering operators were turned on and no initial population members were provided. The XCS was run for 30 runs, with each run consisting of 5000 exploitation trials (10000 trials in total). Figure 9 plots the average System Relative Error from 30 runs illustrating that the System Relative Error measure rapidly falls as the accurate optimally general set of classifiers [O] is discovered and takes over the population.

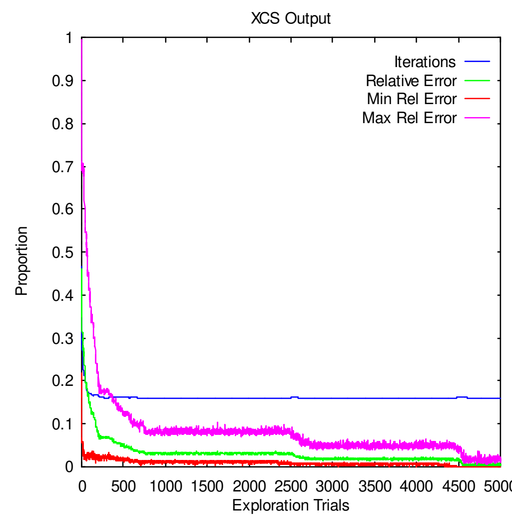

Figure 9 - Fall in error within FSW-9(2) with rule induction

The experiment was repeated within the equivalent FSW with four aliasing states *FSW-9A(2)-4*, with 30 runs and the number of explorations within each run set to 15000 to allow the XCS the opportunity to discover the classifiers before the aliased states. Figure 10 shows the averaged results for the first 5000 exploration iterations and illustrates that, as expected, the aliasing states prevent the System Relative Error from reducing. It is important to note that the System Relative Error values typically 'jittered' around a mean by 0.025 and the Max/Min by 0.07, so the averaging in both experiments has flattened the results.

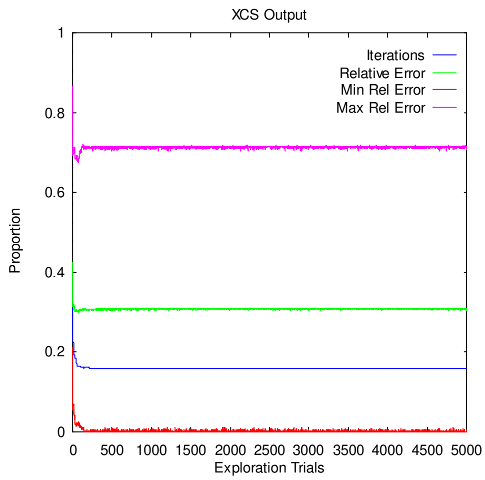

Figure 10 - System Relative Error is not reduced in FSW-9A(2), demonstrating the effects of the aliasing states.

Examining the classifiers produced by the run revealed some very unexpected results. All 30 runs found the following classifiers (taken from one typical run) with high numerosity :

Classifier	Predict.	Error	Fitness	Acc.	N.	MS	Exp.
##1##->0	1000.00	0.000	1.000	1.000	74	75	30000
##1##->1	710.000	0.000	1.000	1.000	71	75	14901
##0##->0	294.104	0.189	0.679	0.000	50	79	209457
##0##->1	138.055	0.004	0.933	1.000	87	97	104850
###10->0	171.518	0.000	0.990	1.000	34	108	29593
###0#->0	147.872	0.002	0.968	1.000	34	101	59824

All the remaining classifiers were not uniformly represented across the populations, and had low numerosity (mean N<3). Two co-acting reasons can be identified for this. Firstly, the only competing classifiers in [M] for these states will be less general but no more accurate, more general and less accurate, or of the same generality but selecting a lower rewarding action. Thus, the competition within the match set is insufficient to put deletion pressure on the classifier. Secondly, and more substantially, the hypothesis failed to account for the fact that in an environment with consecutive aliasing states XCS will dwell in the aliased states for proportionately longer than the other states and will therefore provide more opportunities for the GA to be invoked. Less general classifiers will be no more accurate and will compete for GA involvement less often and so will be eradicated, whilst more general classifiers will have a lower accuracy and therefore a lower fitness and be selected for GA use less often. Thus the classifier covering the aliasing states will put deletion pressure on competing classifiers. Furthermore, the increased frequency of GA invocation negates any potential deletion pressure from classifiers covering other match set niches. As a result, the classifier covering the aliasing states is maintained within a population niche inspite of its inaccuracy, contrary to Hypothesis 3.

A further unexpected result was that the classifier covering s_2 action 0 was present, but the classifiers covering s_0 and s_1 action 0 were represented by an over-general classifier and no high numerosity classifiers covered action 1 in these states. In order to identify a reason for this result, the experiment was re-run and the predictions of the expected members of [O] were recorded. An examination of the match sets and predictions demonstrated that the aliasing classifiers not only covered the aliasing states but also the preceding states, thereby competing with these classifiers in each GA. It was hypothesized that, whilst this will cause the prediction of the classifier covering the aliased states to oscillate more, since it is already inaccurate this has little effect compared to the benefit of being involved in more match sets. The classifiers covering s_0 to s_2 have a higher prediction as a result of the payoff directly from the classifier covering the aliasing states, and the prediction of s_0 and s_1 is sufficiently close to replace them both with one more general classifier. An experiment that attempts to verify this finding by contradiction was constructed. The stimulus presented by each state was changed so that the aliasing states would provide a stimulus sufficiently different from all other stimuli so as to make the appearance of the classifier covering the aliased states in the match sets of other states improbable. If the hypothesis is correct, the other classifiers should be able to form stable prediction and high numerosity values without disruption from the inaccurate classifier. This experiment gave stimuli to states as follows: 0- 11110, 1- 11101, 2- 11011, 3..6- 00000, 7- 10111, 8- 01111. The experiment consisted of 30 runs of the *FSW-9A(2)-4* environment, and the resulting populations were analyzed. In all 30 runs all of the State × Action × Payoff mapping was covered demonstrating clearly that the classifier covering the aliasing states previously was interfering in the match sets and preventing the formation of accurate competing classifiers.

Given the results obtained, is it possible to set aside Hypothesis 3? Unfortunately, the experiments neither affirm nor deny the hypothesis because the tests failed to provide the classifiers in the aliased states with credible competition but rather afforded them more GA opportunity which instead encourages the maintenance of the inaccurate classifiers. Therefore, a modified FSW was constructed based upon *FSW-9A-2*. The two alias state test environment was chosen to minimize disruption to preceding classifiers, reduce the prevalence of the aliased states within a GA, and thereby increase competition. To further increase competition the environment was extended to provide four actions in each state. One action ('00') moved the FSW into the following state, and all other actions kept the FSW in the current state. The increase in actions increases competition for population space further. To prevent the aliasing states looping between themselves an additional state was provided. This state was obtainable from all the aliasing states by transitions labeled with the actions 01, 10, and 11. As for the other non aliased states the only transition out of the state was labeled with action 00 and moved to s_7. Finally a message encoding for the states was chosen to prevent the classifier covering the aliasing states interfering with preceding classifiers. The following encoding, selected from three devised encoding attempts, produced the minimum interference: 0- 00111, 1- 00110, 2- 00011, 3- 0010, 4- 00100, 5..6- 11111, 7- 00000, 8- 00010, 9- 00001. XCS was run ten times within this environment using the same parameters as the previous tests.

Table 1: The optimal classifiers produced within the modified *FSW-9A(2)-2* showing the failure to establish long-term classifiers covering the aliased states.

Classifier	Mean Numerosity	Mean Match Set	Mean Experience
##010->00	13.5	16	29043
##010->01	10.8	12.9	14229
##010->10	11.2	13.7	13315
##010->11	9.8	11.9	14238
##00#->00	14.4	17.6	29460
##00#->00	10.7	13.4	14601
##00#->00	11.8	13.8	14533
##00#->00	11.1	13.2	14438
1####->00	1.7	10.4	43
1####->01	8.3	10.2	7010
1####->10	8.9	11.1	4834
1####->11	8.5	10.5	7678
##100->00	14.5	17.4	27188
##100->01	11.4	13.3	14578
##100->10	10.8	12.6	12980
##100->11	10.9	12.7	12713
###01->00	13.7	17.5	24960
###01->01	12.7	14.8	12991
###01->10	10.8	14.4	13972
###01->11	10.1	12.2	13503
##0#1->00	12.5	16.2	27671
##0#1->01	10.6	13.3	15874
##0#1->10	10.9	14.0	14524
##0#1->11	9.8	12.4	14741
##110->00	13.5	16.5	25447
##110->01	11	12.5	10760
##110->10	11.1	13.1	13801
##110->11	9.5	11.72	12962
0#111->00	14.3	17.2	21601
0#111->01	11.2	13.1	11658
0#111->10	9.2	11.7	15090
0#111->11	10.9	12.7	12312

The resultant populations were captured and examined to identify [O] and look for evidence of the competition driving the classifiers covering the aliased states out in the search for population coverage. Table 1 lists the contents of [O] with averaged numerosity, match set membership, and experience. In all runs [O] was found to be completely represented - there was no interference from the classifier covering the aliased states. In all populations the rule representing forward movement within the aliased states was provided by a number of low experience (mean experience 50, compared with [O] mean of 15893) classifiers with low numerosity (mean numerosity 1.7, compared with [O] mean of 11.2) but with a Match Set estimate equivalent to other members of [O]. This demonstrates that the classifiers representing the aliasing states are in competition between themselves to find an accurate generalization but because of the pressure for population space no one inaccurate classifier is able to dominate the population niche and the XCS continues in perpetual ineffective exploration for a suitable accurate classifier, as Hypothesis 3 predicted.

4 DISCUSSION

The work of Lanzi is the only other commentary on the learning of aliased states within XCS. The work which has been presented in this paper is complementary to that work, verifying the phenomenon, and identifying the effect that aliased states will have on the attempt by XCS to establish an accurate, optimally general and complete State × Action × Payoff mapping. Important additions to the current body of work are the identification of the Consecutive State Problem, which Barry (1999) goes on to demonstrate is a sub-problem of the Aliasing Problem, the establishment of the disruption caused to existing classifiers covering the preceding states, the demonstration of extensive disruption caused by additional match set involvement by over general classifiers covering the aliased states, and the establishment of the conditions under which classifiers covering the aliased states are established or eradicated by the action of the induction operators.

Lanzi (1997) reported that within tests using the XCSM1 memory solution to the Aliasing Problem within the environment Woods102 the modified XCS failed to obtain a reasonable performance whenever non-uniform exploration rates were used. "Most important the system fails to converge to optimal performance when, due to the structure of the environment, the agent is not able to visit all the areas of the environment uniformly" (Lanzi, 1997, section 7.3). The results obtained from the non-uniform run without learning in the simple FSW-5A environment, depicted in figure 7, shed some light on Lanzi's findings. Non-uniform exploration will cause the prediction of the classifiers to fluctuate much more than under uniform exploration; therefore the disruption to both classifiers attempting to cover the aliased states and those covering preceding states will be more significantly affected.

On the application of XCSM to the more difficult MAZE7 environment Lanzi (1997) reported difficulties in establishing effective performance until exploitation only runs were used. Given the distance between the aliased states in MAZE7, the oscillation in prediction of any classifier attempting to cover these states will be large. From the results presented in this paper, it is possible to hypothesize that the large oscillation will not only impact the preceding classifiers, but may also encourage the formation of one or more over-general classifiers which could be involved in the action sets of two groups of three states from the nine states within MAZE7 and thereby adversely affect the formation of adequate classifiers for the remaining intervening states. Certainly, the inability of the memory mechanism to successfully disambiguate the two aliased positions only under combined exploration/exploitation would suggest some degree of prediction inaccuracy, and would support Lanzi's choice of a hybrid mechanism utilizing exploitation only for the memory mechanism (Lanzi, 1998). Unfortunately Lanzi does not provide any error measures or population details that would allow more specific hypotheses to be generated without repeating his experiments.

5 CONCLUSIONS

This work has contributed to XCS research in a number of ways. Firstly it has provided replication of the problems arising from the presence of aliased states within an environment, a phenomenon first identified by Lanzi (1997). Secondly, it has clarified the phenomena identified by Lanzi by simplifying the environment used, and demonstrating that this phenomena can occur as the Consecutive State Problem, which Barry (1999) demonstrates is a sub-problem of the Aliasing Problem. This result is important to many problem domains. For example, within robotic control, corridor following behavior will be an instance of the Consecutive State Problem but clearly corridor following behavior is a fundamental robot skill. Thirdly it identified the extent by which the aliasing inaccuracy can affect classifiers representing preceding states within the environment, and discovered (fourthly) that classifiers covering the aliased states can even take over other [M] in order to gain advantage within the GA but to the detriment of classifiers covering non aliased states. Fifthly it has shown that the intuitive conclusion that the classifiers covering the aliased states would be unable to sustain themselves within the XCS population is unfounded except in those cases where strong competition for GA involvement and population space prevent an inaccurate classifier dominating a population niche purely on account of an accidental increase in its numerosity. However, irrespective of the ability of classifiers covering consecutive aliased states to sustain themselves, the existence of the aliased states will prevent XCS from finding and maintaining [O] and therefore will impact the decision making ability of XCS in more complex environments.

Further work is required in order to establish the application of these results to more complex environments. Further work is also required to compare the results produced with effects noted by Lanzi (1997) within the more general Aliasing Problem. In particular, the duplication within the MAZE7 environment (an example of separate state aliasing) of the results on the disruption of the population by aliasing classifiers due to additional match set involvement would lead to a greater understanding of the effects on learning over non-aliased states by the attempts to learn coverage of the aliased states within all aliasing environments.

Acknowledgements

The author wishes to acknowledge the discussions and advice from Stewart Wilson and Tim Kovacs in the development and testing of the XCS implementation used in this research, which is freely available from:

http://www.csm.uwe.ac.uk/~ambarry/LCSWEB

References

Barry, A.M. (1999), Aliasing in XCS and the Consecutive State Problem : 2 - Solutions, submitted to the Intl Conf on Genetic and Evolutionary Computing, 14-17 Jul 1999.

Booker, L.B. (1982), *Intelligent behaviour as an adaptation to the task environment*, Ph.D. Dissertation (Computer & Communication Sciences), Univ. Michigan.

Cliff, D., Ross, S. (1994), Adding memory to ZCS, *Adaptive Behaviour 3(2)*, 101-150.

Grefenstette, J.J. (1987), Multilevel Credit Assignment in a Genetic Learning System, in *Proc. Second Intl. Conf. On Genetic Algorithms and t heir Applications*, 202-209.

Kovacs, T. (1996), Evolving optimal populations with XCS classifier systems. Tech. Rep. CSR-96-17, School of Computer Science, University of Birmingham, UK.

Kovacs, T. (1997), XCS Classifier System Reliably Evolves Accurate, Complete, and Minimal Representations for Boolean Functions, WSC2: 2nd On-line World Conference on Soft Computing in Engineering Design and Manufacturing, 23-27 June 1997.

Holland, J.H. (1986), Escaping Brittleness : The possibilities of general-purpose learning algorithms applied to parallel rule-based systems, in M. Michalski, M. Carbonell (Eds), *Machine Learning, an Artificial Intelligence approach, Vol. 2*, pp. 593-623.

Lanzi, P.L. (1997), Solving problems in partially observable environments with classifier systems, Tech. Rep. 97.45, Dipartimento di Elettronica e Informazione, Politecnico do Milano, IT.

Lanzi, P.L. (1998), An analysis of the Memory Mechanism of XCSM, in *Proc. Intl Conf. on Genetic Programming.*

Riolo, R.L. (1987), Bucket Brigade performance: I. Long sequences of classifiers, in *Proc. Second Intl. Conf. on Genetic Algorithms and their Applications*, 184-195.

Wilson, S.W. (1983), Knowledge growth in an artificial animal, in *Proc. First Intl. Conf. on Genetic Algorithms and their Applications*, 196-201.

Wilson, S.W. (1994), ZCS, a zeroth level classifier system, *Evolutionary Computation 1(2)*, 1-18

Wilson, S.W. (1995), Classifier fitness based on accuracy, *Evolutionary Computation 3(2)*, 149-175

Wilson, S.W. (1998), Generalization in the XCS classifier system, in *Proc. Third Annual Genetic Programming Conference* (GP-98).