Evolving Tactical Formations on the RoboCup Field

Michael de Raadt[1], Mikhail Prokopenko[2], Marc Butler[2]

[1] University of Southern Queensland
Department of Mathematics and Computing
Toowoomba, Queensland, 4350, Australia
deraadt@usq.edu.au

[2] CSIRO Mathematical and Information Sciences
Intelligent Interactive Technology
Locked Bag 17, North Ryde, NSW 1670 Australia
{mikhail.prokopenko, marc.butler}@csiro.au

Abstract. Evolutionary methods are well recognised as having the powerful potential to adapt agents in multi-agent complex systems. This paper describes an experiment where a variant of typical evolutionary methods (genetic algorithms) was used to evolve a family of agents in such an environment. The RoboCup environment is an excellent example of a complex system where individual agents must process fragmented, localised and imprecise information, and act autonomously, while achieving a reasonable degree of cooperation. The principal aim of this study was to examine the use of evolutionary methods in adapting agent's properties, instead of hard coding and iteratively fine-tuning approximated parameters. Tactical formations of a simulated soccer team were chosen as a subject for experimentation. Preliminary empirical results showed that distinct new tactical formations emerged, exhibiting a degree of adaptation not only to the environment but also to the underlying agent architectures.

1 Introduction

The Robot World Cup Initiative (RoboCup) is an attempt to foster research in various technological fields by providing a standard problem where a wide range of technologies can be integrated and examined. For this purpose, RoboCup chose to use the game of soccer, and organized RoboCup: The Robot World Cup Soccer Games and Conferences. RoboCup provides a forum for building advanced physical and synthetic agents that can play soccer at an increasing level. The RoboCup Simulator is a soccer simulation system which allows virtual soccer players to compete in a dynamic, complex, uncertain multi agent environment. The soccer server attempts to provide a challenging environment for Artificial Intelligence research, by allowing researchers to concentrate on designing "brains" for the simulated bodies of their players. The simulation system is run as a client/server system, with a central simulator server communicating with 22 player clients [3]. Each client receives sensory data (visual, auditory, internal) about its own local view of the world (from its own position), and can send commands such as run, turn, kick, etc. As in the real game, "players" have only fragmented, localised and imprecise information about the field, and only have a short time to respond. They must make autonomous decisions and act in collaboration with members of their own team but in opposition to another team, with the aim of scoring as many goals as possible, while defending their own goal.

An experiment was conducted which aimed to improve a team which competed successfully in previous RoboCup competitions (Simulation League). This team of autonomous heterogeneous software agents was designed and implemented using a hierarchical logic-based framework for synthetic agent architectures. The framework is expressive enough to capture a subset of desirable properties from both the situated automata and subsumption-style architectures, while retaining the rigour and clarity of logic-based possible worlds semantics [5, 6, 7, 8]. Basic skills for each agent were mostly hard-coded. The aim of this paper is to demonstrate how evolutionary methods can be used to evolve an agent's properties, instead of hard coding and iteratively fine-tuning approximated parameters. In particular, tactical formations were chosen as a subject for experimentation.

Defining behaviors through evolutionary methods is not novel. Genetic Algorithms (GA) have been used in the past to develop a range of behaviours from low levels of functionality like moving a robot backwards and forwards [9] through to coordination in multi-agent environments [10]. In the world of RoboCup, Genetic Programming has been used to evolve behaviors of simulated soccer players with some degree of coordination [4].

2 Motivation

Each player has a set of "properties" that holds values defining how that player responds to the environment while playing a game. Examples of properties include how close a ball is before the player will attempt to chase it, how fast the player will run when the ball is a certain distance away from it, etc. Each player in a team has properties in common with other players and properties unique to its tactical role (e.g., fullback). It is our intention to extract these properties and represent them in a structure uniform across an entire team (even though not all properties are immediately relevant for each player). Each player will draw from a single set of properties that in turn defines the team behaviour. Let us call these properties a player's genome. In this sense the players are homogenous (share the same genome).

Keeping the properties of each player in a single structure does not remove their individuality. But a team should not be thought of as a population of individuals, instead, it should be considered as a family of individuals that evolve together within a population of families. Representing the properties of a whole team in a single structure has advantages and disadvantages. The main disadvantage is redundancy of information - each player carries the parameters required for all tactical positions. However, this could also be seen as an advantage. When only part of the genome changes through evolution, the remainder is inherited between generations. This keeps non-relevant genes "passive" but technically accessible to other team-mates (family members). Also, with the genes of other players available, two players may dynamically re-activate passive genes, and essentially swap tactical positions during a game if required.

Historically, the game of soccer has developed along a tactical dimension, where new formations were continually required to "ensure an ever more effective distribution of the eleven soccer players on the field" [2]. The experiment aimed to explore which tactical formations (family configurations) better suited the basic skills of a given team, having basic agent architectures fixed for the duration of the experiment. In other words, the goal of the experiment described here was to produce the ultimate team (family), which may not necessarily be based on individual brilliance, but rather accommodated all players in their most efficient capacities.

By coercing the teams to evolve in constrained ways, it was hoped that behaviours would emerge that will advance the fitness of a team. Such emergent behaviours may be the result of minor changes to properties (mutations) that may have larger overall effect on the team as a whole. By combining parts of successful teams (through crossover) it may be possible to capitalise on their fitness and produce better teams. Such changes may be counter-intuitive (i.e. it may be thought that such changes may have a negative effect on the players), but in context of team-work, such changes may prove beneficial.

3 Experiment

3.1 Tournament Selection

Each team's genome was stored in a file that each player draws from when it is initialised. The players are then sent into a game on the RoboCup field.

One obstacle was deciding which team should be the opposition for a team to prove its fitness. It was conceived that six teams would be created through random or deliberate mutations of an initial team. Each team was distinguished by its own genome. Then teams played each other and results were recorded. The results of a 'round-robin' tournament were passed to the Evolver (see Fig 1), for subsequent analysis and determination of fitness for each team.

The Evolver consists of several parts. First the fitness of each team was determined. Then, using this information, the next generation of teams was created. The Evolver was also used to maintain a historical record of each generation for future analysis, or for use if evolution seems to head down a wrong track or stagnates.

Fitter teams were used as the basis for each "next" generation (created by the Team Generator). Once teams had played they were ordered. Less fit teams were lost and fitter teams survived to reproduce, effectively leading to evolution. Fitness was based on raw score with account to the difference of goals scored. Analysis of other factors, like time in possession etc., which may have encouraged the emergence of undesirable behaviours in teams were not used as this multi-factored fitness assessment has been shown to lead to behaviours that may be considered sub-optimal [4].

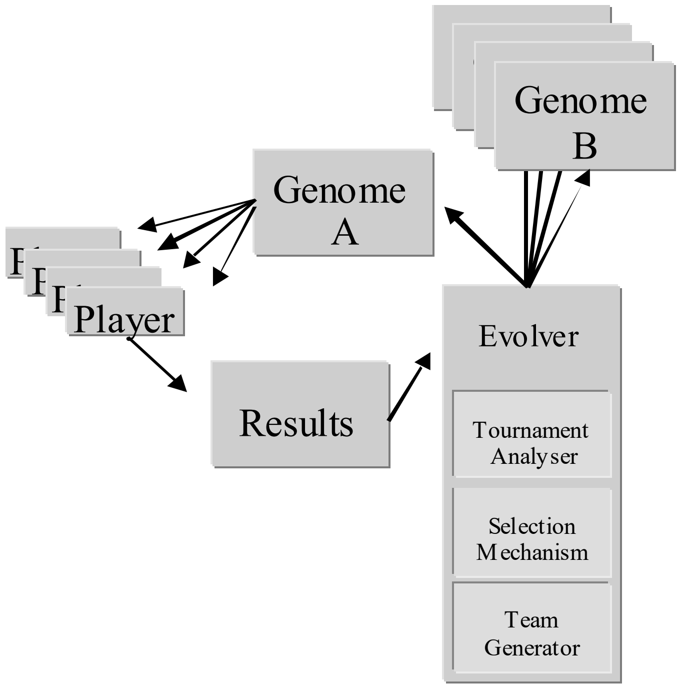

Fig. 1. The Evolver and its part in evolving new team genomes.

One undesirable by-product of synthetic evolution is the tendency for evolving agents to develop skills leading to superior fitness over others in the population, while neglecting other skill improvement. One method used to avoid this was to take successful teams from previous generations and reintroduce them into new generations. This ensured an even development of skills. Successful teams also tested against other teams from the world of RoboCup to determine how successful they are and if the emphasis of our evolution needed to be changed. Even without such adjustments, the teams surviving through the experiment were expected to exhibit property improvements appropriate to their given agent architecture.

3.2 Evolution

Evolutionary techniques used in the experiment included typical reproduction, mutation, and crossover. Reproduction was implemented by taking a successful genome from one generation and reproducing it in the next. Mutation of a team involved taking a (successful) team and changing some part of its genome, then placing it in the next generation. The evolution that was fostered by the Evolver was not purely random. A mutation domain constrained each property so that unfeasible mutations (e.g., locations outside the soccer field or outside respective zones) were not accepted. The evolution was, therefore, pointed towards what was loosely seen as the right direction to speed up the experiment.

Crossover was implemented by taking part of one successful team's genome and inserting it into another successful team's genome to produce offspring for the next generation. Teams swapped equivalent parts of their genome to produce offspring for the next generation. Again, due to time constraints, this was constrained by the Evolver so only parallel parts of two genomes were crossed-over (as opposed to crossovers that exchange sections of the genome that overlap such boundaries).

4 Method

In preliminary runs of the experiment it was decided to manipulate only the zone and location parameters for the team's players. These parameters relate to where players will situate themselves when they are attacking, defending or at a kick-off, and what zones they will cover in offense and defense. It was felt that this would not only achieve our main target in evolving tactical formations best suited for the team, but will also allow us to observe results in a very obvious way. An evolution viewer allowed viewing of changes in locations and zones, and tracking of the heredity of high ranking teams.

Not surprisingly, a major hindrance to this experiment was time. Alterations were made to the Soccer Server to allow playing of shorter automated games (running without human intervention). This was done to allow the experiment to run by itself for extended periods for several generations.

4.1 History of Evolution

The experiment was controlled by a script which ran iteratively through each generation. With n teams there is $n(n-1)/2$ games in a round robin tournament. With six teams, fifteen games need to be played. Each game's score is recorded in the scores file under a label which identifies its generation. At the end of the tournament, the Tournament Analyser (part of the Evolver) created a tournament table which ranked the teams according to their success. In each game a win was awarded 3 points, a loss 0 and a draw gave both teams 1 point. Goals scored for and against each team are used in tie-breaks. A fragment of the scores files is depicted in Fig 2.

The tag "g 3" denotes the beginning of the generation three. This is followed by the information about the first game of the tournament - a game between team 1 and team 2 with a final score of nil all. The second game was between teams 1 and 3 with a score nil-one and so on. After the final game, these scores are read by the Tournament Analyser, which ranks teams according to their success. The ranking table is then used to determine how teams will be evolved.

As well as serving as a basis for judging fitness this scores file was also kept as a log of how the evolution progressed, being essentially a history of the evolution. Together with genome files from various generations, the scores file was used to show paths of evolution and other patterns.

```
Rank: 6 Team: 4 Points: 3 For: 0 Against: 2
------------------------------------
g 3
1 2 0:0
1 3 0:1
1 4 1:0
1 5 0:0
1 6 0:1
2 3 0:0
2 4 0:0
2 5 0:0
2 6 0:0
3 4 0:0
3 5 0:1
3 6 0:0
4 5 0:0
4 6 0:1
5 6 0:0
--- Result table for Generation 3 ---
Rank: 1 Team: 6 Points: 9 For: 2 Against: 0
Rank: 2 Team: 5 Points: 7 For: 1 Against: 0
Rank: 3 Team: 3 Points: 6 For: 1 Against: 1
Rank: 4 Team: 2 Points: 5 For: 0 Against: 0
Rank: 5 Team: 1 Points: 5 For: 1 Against: 2
Rank: 6 Team: 4 Points: 3 For: 0 Against: 2
------------------------------------
g 4
1 2 0:0
1 3 0:1
```

Fig. 2. Results of a generation of evolution and the ranking of team success.

4.2 Next Generation

There are a few selection schemes frequently used in GA: roulette wheel cost, roulette wheel rank and tournament selection. According to the roulette wheel method based on cost function, individual's contribution to next generation is directly proportional to the value of some cost function. This method, however, may have problems when the fitness values differ significantly. When the best fitness value is close to 100%, other individuals will have very few chances to be selected. The roulette rank selection scheme ranks the population from 1 to N, and then assigns fitness values based on the reverse ranking, so that the highest-ranked individual receives N, and the lowest-ranked one is assigned 1. Both roulette wheel methods are probabilistic, and the fittest individual may still miss a chance to contribute to the next generation. On the contrary, an elitist genetic algorithm (like tournament selection) always retains in the population the best individual found so far.

Our task was complicated because our tournament had a small number of teams, and in addition, we varied only a small number of parameters. Therefore, there was a high risk of loosing good solutions with roulette wheel methods, and at the same time, a pure elitist scheme would have been limited by the size of population. A hybrid method was used – it has elements of all 3 selection methods: ranking, the elitist scheme and the cost-based selection. Intuitively, after ranking the teams at the end of each round-robin, we needed to retain the fittest team *without* any alterations, and refuse the two worst performing teams any chance of surviving. This restricted the cross-over operations to (n-2) teams (4 in our case), while there were only (n-3) places to be filled. The natural solution was a cascade crossover: where the first team could be crossed-over with the second, second with the third, and so on, resulting in (n-3) new teams (3 in our case). This controlled cascade crossover rewards the very first team even more, and takes away from the last accepted team on (n-2)-th rank. Mutations were limited only to the last 2 teams – produced from the first two.

Table 1 shows the scheme used to create six new teams for the "next" generation. Let R(x) denote the reproduction of a genome x, M(x) a mutation, and C(x,y) a crossover of x and y.

Table 1. How teams in the next generation a created.

Team in next generation previous generation	Composed from
1'	R(1)
2'	C(1,2)
3'	C(2,3)
4'	C(3,4)
5'	M(1)
6'	M(2)

The six new teams were created based on the results obtained from the previous generation. In order to show that the Evolver rewarded teams proportionally to their success, we can try to approximate the cost function used implicitly in the *cascade crossover*. Let us scale the rewards where the unaltered reproduction receives highest score, and can be given a value of 1. Mutation scores almost as high with most of a team's genetic information carrying into the next generation, so we will value mutations at 1 also. Crossover projects, on average, half of a team's genome into the next generation so we value each crossover as ½. Hence, the cost function can be approximated[1] as follows (table 2):

Table 2. Approximated cost function.

Rank	Reward	Reward Points
1	Reproduction (1 point), Mutate (1), Crossover with 2 (½)	2 ½
2	Mutate (1), Crossover with 1 (½), Crossover with 3 (½)	2
3	Crossover with 2 (½), Crossover with 4 (½)	1
4	Crossover with 3 (½)	½
5	Did not survive	0
6	Did not survive	0

[1] Importantly, this cost function approximation is only an underlying justification – it is not used in the cascade crossover selection scheme.

5 Preliminary Results

The experiment began using a formation as appears in Fig 3, consisting of 3 defenders, 5 midfielders and 2 attackers in a symmetric pattern. During a game, every player was capable of distinguishing (with a reasonable degree of success) which team, own or opponent, has possession of the ball. Dependent on this information, the player tries to situate itself in an attacking or defensive location, respectively. It is worth noting, however, that being at an attacking location does not necessarily mean that the player is going to support attack through to the opponent goal, but rather that it tries to be useful in a context of team-work. For example, possessing the ball may require team-mates to be in open locations ready to receive a direct pass, or making space for medium-range dribbling. Similarly, the defensive location may appear to be quite exposed but the player may find it beneficial to exploit an off-side trap[2]. The initial "symmetric" tactical configuration for attacking and defending locations were the same, but evolved independently.

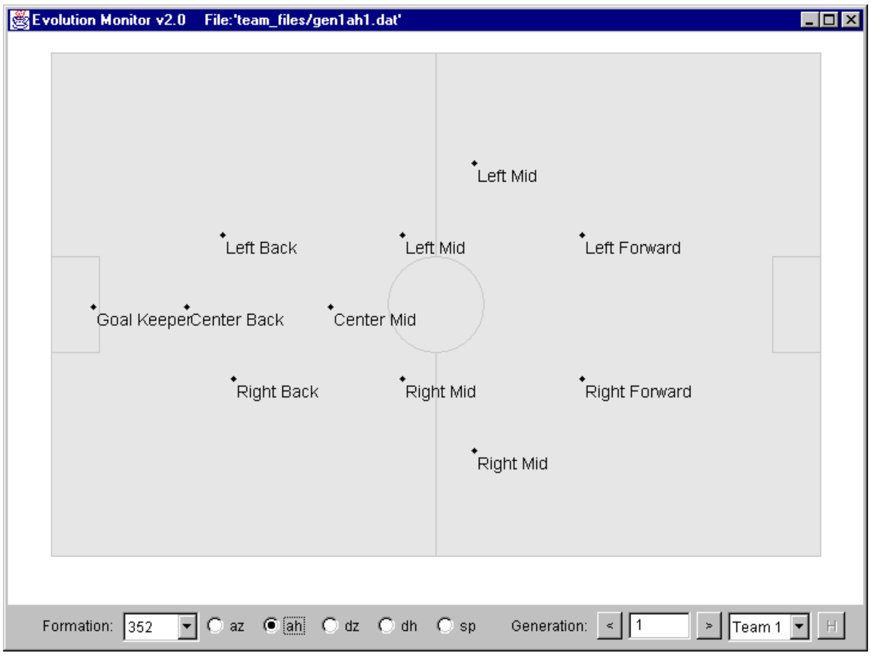

Fig. 3. Attacking and defensive locations of one team at start of experiment.

The experiment was run for 256 generations, with 3840 games played. First observable trends showed periods of intense and roughly equal competition, followed

---

[2] An off-side is called when a pass is directed towards players who have less than two opponents between them and the opponent goal at the moment of the pass – the rule prevents all attacking players concentrating right in front of the opponent goal, waiting for a long pass. Defenders may move forward creating an "off-side trap", which forces opponent attackers to back up as well.

by the emergence of a mutant that would dominate for a few generations (up to 5 – 6). After another few generations, other teams appeared to 'catch up' and absorb the dominance of the leading team. This was due to the fittest team's genetic advances propagating into other genomes.

The first perceptibly distinct formations appeared at around generation 70 (Fig 4). The fittest teams showed attacking locations that did not appear to be very efficient. While the wing midfielders have moved out to the wings and two forwards basically kept their locations, all three backs and three midfielders clustered in the center. Obviously, such a formation does not tend to be very aggressive: a ball cannot be swiftly and efficiently delivered from deep defense forward. Not surprisingly, not very many goals were scored at this time, with most games ending in a nil all draw, there were on average only 5-8 goals scored in each generation (15 games). However, on analysis, these formations could be seen as a result of behavioral rules that were fixed in the architecture of each player. It is important to mention, at this stage, that our players have a tendency to pass a ball when it is in the own half, and dribble the ball when it is in the opponent half. Although, this is only a general rule, and varies dependent on a context, it can partially explain a visible gap between the mid-field and the attack. Nevertheless, such a tactical formation (we called it the "cluster" formation) was slightly disappointing.

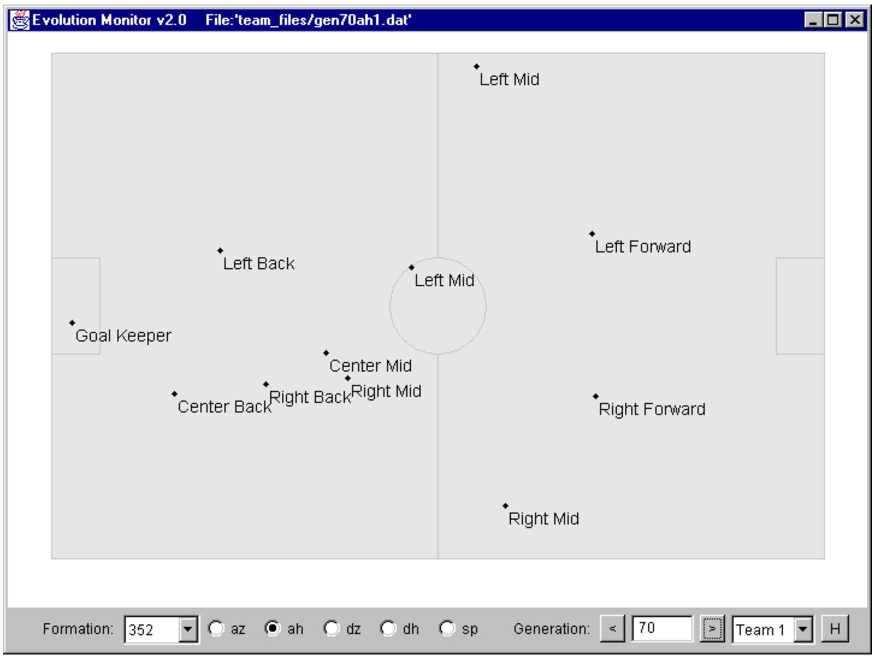

Fig. 4. At around generation 70, the fittest team showed attacking locations as above. Note that the wing midfielders have moved out to the wings, while backs and centre-midfielders cluster in the center.

With evolution progressing, certain new formations appeared. It took in fact about 150 generations to evolve a substantially new formation, observed from early 210's generations till the end stage. Fig 5 shows generation 256's highest ranked team's attacking home locations. There are still two forwards close to the opponent's goal, but not too close to be trapped off-side (off-side rules have very strong influence in the RoboCup world). There are two midfielders at the wings with plenty of room to move forward. Wing runs are quite suitable for our players given their dribbling skills. With no players within range to pass to, players (like the wing right mid-fielder) will carry the ball forward while dribbling and keeping possession of the ball, effectively ensuring that the team stays in attack. The remaining players, including the goalie, appear to be evenly distributed in a formation that covers an area from the bottom-left to the top-center of the field. This is definitely due to the passing behavior that the players follow: when a team-mate is within a passing range and not too close to a player with the ball, a pass will be attempted.

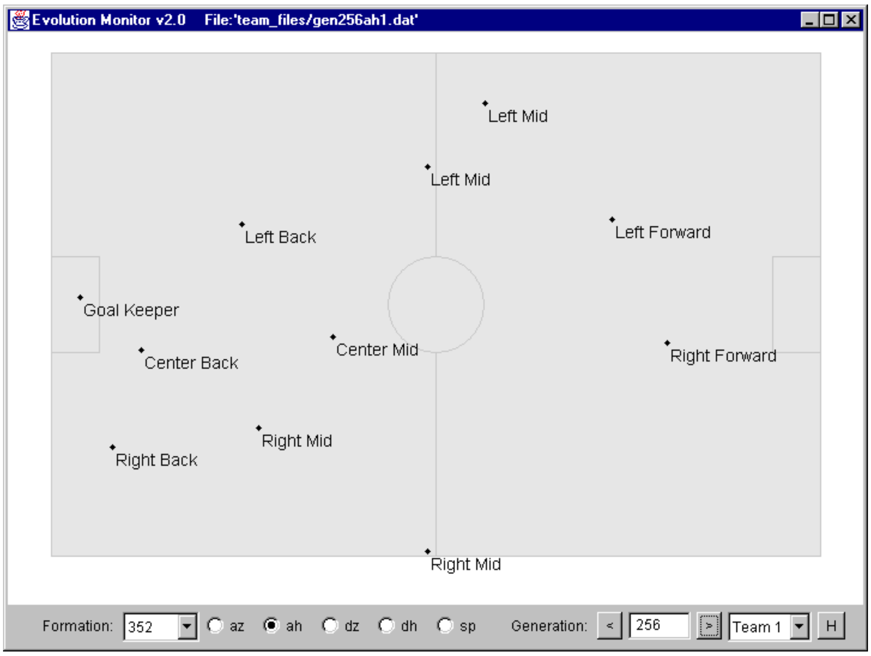

Fig. 5. At generation 256 we see an attacking formation with forwards close to the opponent's goal, backs and some midfielders dispersed equidistantly and 2 wing midfielders with plenty of room in front of them.

With players distributed as in Fig 5, there are good candidates for passing and such a distribution can quickly carry the ball out of their own half and to the forwards. Besides, the forwards themselves are not too isolated – a good 20 meters[3] pass from the left wing mid-fielder to the left forward is likely to succeed. This formation (we called it the "ladder" formation) led rather naturally, to higher scoring results (10 – 12 goals per tournament, on average).

[3] The size of the simulated soccer field is approximately 105 x 68 meters.

It was quite unexpected to observe an asymmetric formation emerging (and surviving) after 256 generations, given that agent behaviours are totally left-right symmetric. However, these empirical results show again that evolution may take strange ways, and overall emergent behavior of the fittest teams may not be reduced to a sum of components [1].

In Figs 6 and 7 we see the evolution of the defensive formation. The defensive formations of the fittest teams at around the generation 70 did not bring any surprises, beyond the observation that the goal-keeper has moved quite far forward (in fact, the goalie is more advanced than the centre-back). Importantly, when an opponent has the ball, the team may choose to cover all areas of the field, and the observed defensive formation copes with this quite well resembling the “balanced” formation, well recognised in soccer tactics. It is interesting that this formation is quite effective against the “cluster” formation observed in attacking formations at this time - in particular, completely shutting down chances of centre penetration, and limiting wing runs. This, quite likely, explains low scores of early generations as well.

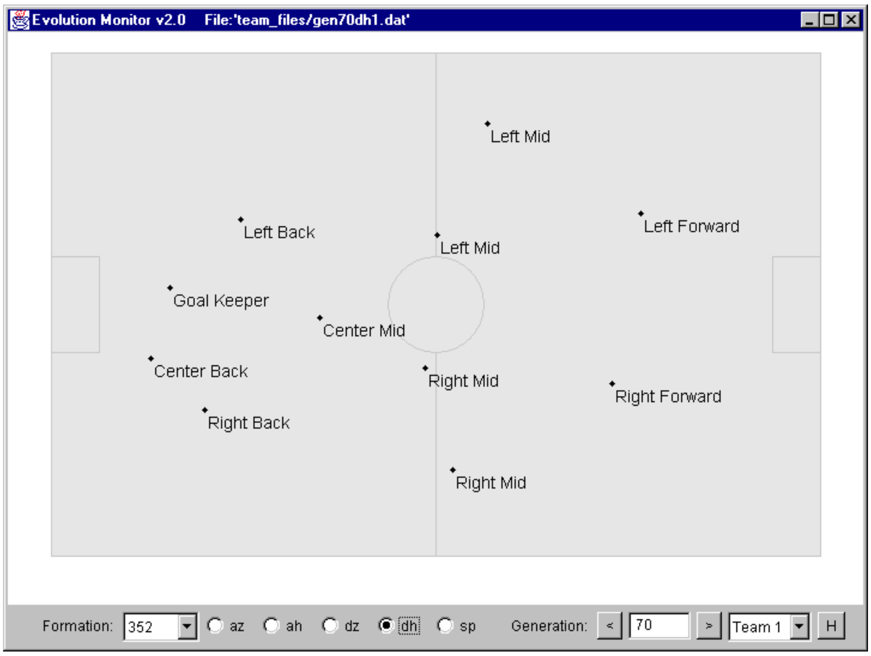

Fig. 6. At generation 70 not much has changed in defending locations from generation 1. Note though that the goalkeeper has come forward to be among the backs.

At generation 256 pictured in Fig 7, we can observe that the goalie and other backs have moved forward, risking an empty goal. The goal-keeper practically abandoned its place, and the center-back mutated to cover the goalie area to the best of its ability. It should be noted that the test-goalie skills are not world-class at the moment, so the swap was not as unfair as it seemed (the centre-back ball chasing skills often did as good job as goalie interceptions). A risky defensive formation could be attributed to the off-side rules of RoboCup.

These locations do appear to 'trap' opponents in an offside, which results in a turnover of possession and moves the attacking opponent away from the ball. This "opportunistic" formation may explain the higher scores as well. Although the mid-field is covered well, the forwards are quite far advanced, with the left forward also in an opportunistic location, bordering on off-side. It should be remembered that defensive locations are approached when players try to win possession back, and so, although forwards are close to the goal, they are not off-side unless they are being passed the ball.

The analysis showed that, tactically, the "cluster-balanced" formation was marginally better than the original "symmetric" formation, but lost ground to the "ladder-opportunistic" formation during the evolution.

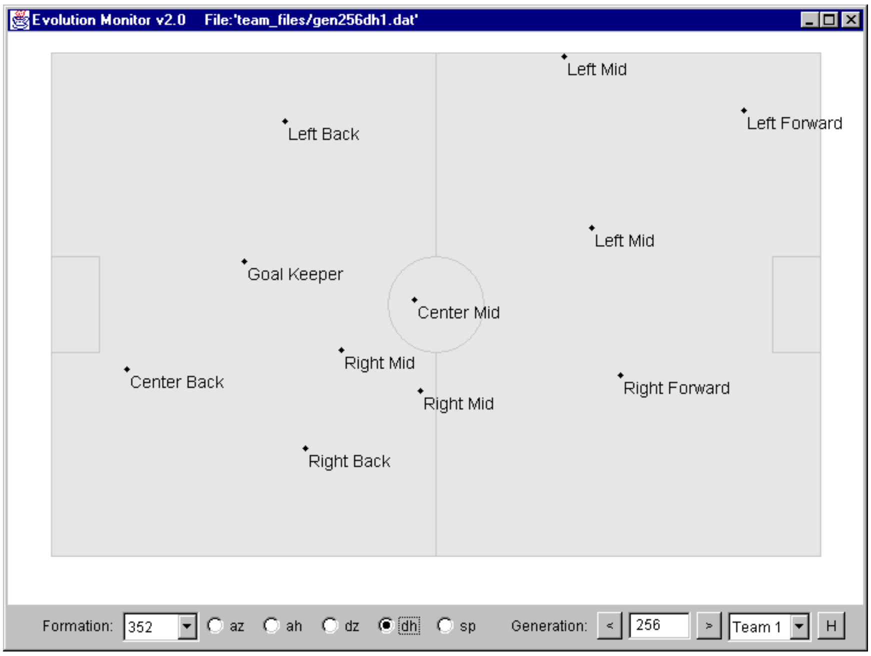

Fig. 7. At generation 256 most players have moved forward. The mid field is covered, the goalie has moved well away from the goal and the forwards are quite far advanced with the left forward in a probable off-side location.

6 Conclusion

Empirical results showed distinct new tactical formations emerging, justifying the cascade crossover method introduced in this work. In addition, a degree of adaptation, not only to the environment but also to the underlying agent architectures, was exhibited.

Further generations and more teams would produce an evolution based less on chance. The experiment was limited to six teams. This could be extended through other tournament styles like a Swiss-style Tournament or by utilising more computing power. Also, the inclusion of one or more static teams (say champions from previous RoboCup tournaments) may result in a more focussed evolution while still preserving a ramped co-evolution of teams. In addition, some teams may be more successful in certain areas like defense or attack rather than in an overall strength – so it is possible to keep such teams in a repertoire of genetic memories by storing their genome. Another future direction is to experiment with agent architectures themselves – now that a framework is in place, it is a matter of extending the players' genome.

References

1. Casti, J. Would-Be Worlds: How Simulation is Changing the Frontiers of Science. John Wiley and Sons. New York, 1997.
2. Gardner P. The Simplest Game. 1996.
3. Kitano, H., Tambe, M., Stone, P., Veloso, M., Coradeschi, S., Osawa, E., Matsubara, H., Noda, I., and Asada, M. The RoboCup Synthetic Agent Challenge. In Proceedings of the 15th International Joint Conference on Artificial Intelligence, Nagoya 1997.
4. Luke, S., Hohn, C., Farris, J., Jackson, G., Hendler, J. (1997), Co-evolving Soccer Softbot Team Coordination with Genetic Programming, In Proceedings of The First International Workshop on Robocop, IJCAI-97, Nagoya, Japan., Springer-Verlag
5. Prokopenko, M., Kowalczyk R., Lee M., Wong, W.-Y. 1998. Designing and Modelling Situated Agents Systematically: Cyberoos'98. In Proceedings of the PRICAI-98 Workshop on RoboCup, 75-89. Singapore 1998.
6. Prokopenko, M., M. Butler. Tactical Reasoning in Synthetic Multi-Agent Systems: a Case Study. In Proceedings of the IJCAI-99 Workshop on Nonmonotonic Reasoning, Action and Change, 57--64, Stockholm, 1999.
7. Prokopenko, M. Situated Reasoning in Multi-Agent Systems. In Working Notes of the AAAI-99 Spring Symposium on Hybrid Systems and AI. Stanford 1999.
8. Prokopenko, M., M. Butler, T. Howard. On Emergence of Scalable Tactical and Strategic Behaviour. In RoboCup-2000: Robot Soccer World Cup IV, Springer, 2000.
9. Steels, L. Emergent functionality through on-line evolution. In Proceedings of Alife IV. Cambridge, MIT Press, 1994.
10. Steels, L. Intelligence - Dynamics and Representations. In Steels, L. (ed.). The Biology and Technology of Intelligent Autonomous Agents. Berlin: Springer-Verlag, 1995.