TEAM: The Team-oriented Evolutionary Adaptability Mechanism

Sander Bakkes, Pieter Spronck and Eric Postma

Universiteit Maastricht, Institute for Knowledge and Agent Technology (IKAT), P.O. Box 616, NL-6200 MD Maastricht, The Netherlands
`Sander.Bakkes@student.tUL.EDU, {p.spronck, postma}@cs.unimaas.nl`

Abstract. Many commercial computer games allow a team of players to match their skills against another team, controlled by humans or by the computer. Most players prefer human opponents, since the artificial intelligence of a computer-controlled team is in general inferior. An adaptive mechanism for team-oriented artificial intelligence would allow computer-controlled opponents to adapt to human player behaviour, thereby providing a means of dealing with weaknesses in the game AI. Current commercial computer games lack challenging adaptive mechanisms. This paper proposes "TEAM", a novel team-oriented adaptive mechanism which is inspired by evolutionary algorithms. The performance of TEAM is evaluated in an experiment involving an actual commercial computer game (the Capture The Flag team-based game mode of the popular commercial computer game Quake III). The experimental results indicate that TEAM succeeds in endowing computer-controlled opponents with successful adaptive performance. We therefore conclude that TEAM can be successfully applied to generate challenging adaptive opponent behaviour in team-oriented commercial computer games.

1 Introduction

In the last twenty years commercial computer games (henceforth called "games") became increasingly realistic with regard to visual and auditory presentation. Unfortunately, artificial intelligence (AI) in games did not reach a high degree of realism yet [1]. In recent years, game developers started to improve AI in their games, focussing specifically on "opponent AI", i.e., the behaviour of computer-controlled agents [2] that compete with a human player.

Opponent AI is typically based on non-adaptive techniques [3]. A major disadvantage of non-adaptive opponent AI is that once a weakness is discovered, nothing stops the human player from exploiting the discovery. The disadvantage can be resolved by endowing opponent AI with adaptive behaviour, i.e., the ability to learn from mistakes. Adaptive behaviour can be imposed on agents by means of using machine-learning techniques, such as artificial neural networks [4] and evolutionary algorithms [5]. In state-of-the-art games, however, adaptive techniques are seldom used.

One area where AI in games can profit from adaptive behaviour is the organisation and interaction of opponents in team-oriented games. Even in state-of-the-art games this so-called "team AI" is barely challenging. The aim of our research is to create more challenging team AI by endowing the opponent AI with adaptive behaviour.

The organisation of this paper is as follows. Section 2 will discuss adaptive team AI in current commercial computer games. The TEAM artificial intelligence adaptation mechanism is discussion in section 3. In section 4, an experiment to test the performance of TEAM is discussed. Section 5 reports our findings, and section 6 concludes and indicates future work.

2 Adaptive team AI in commercial computer games

We define adaptive team AI as the behaviour of a team of adaptive agents that competes with other teams within a game environment. Adaptive team AI consists of four components: (1) the individual agent AI, (2) a means of communication, (3) team organisation, and (4) an adaptive mechanism. We discuss each of these four components below.

2.1 Individual agent AI

Individual agent AI is, as the name implies, the AI of an agent which controls the agent's behaviour. Individual agent AI is game-specific.

2.2 Communication

Coordinated behaviour in a team requires communication. Typically, agents pass along messages containing information or commands. The information can be used to compute counteractions and distribute commands amongst the team-members.

2.3 Organisation

Internal organisation is required to establish team cohesion. Two distinctive approaches to organising a team are: (1) a decentralised approach, and (2) a centralised approach.

The decentralised approach is an extension of the individual agent AI. In the decentralised approach, agents operate in a non-hierarchical communication structure. Figure 1 (left) is an illustration of a decentralised group of agents. Team behaviour emerges through the combined interaction of all agents.

The centralised approach, schematically displayed in figure 1 (right), is strictly hierarchical and is specifically designed to create and maintain well-organised team cohesion. In this approach, the process of decision-making is centralised. Typically, a centralised decision-making mechanism observes agents, makes a decision and processes the decision into agent-commands. The implementation

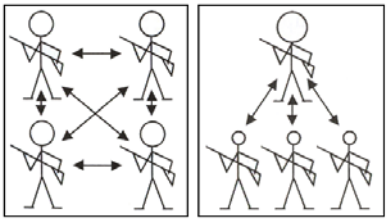

Fig. 1. Decentralised organisation (left) and centralised organisation (right) [6].

of the centralised mechanism varies in style from authoritarian (which focuses on team performance by forcing agents to commands) to coaching (which advises, rather than forces, agents).

2.4 Adaptive mechanism

To our knowledge an adaptive mechanism for team AI does not yet exist in any game. We decided to design and implement an adaptive mechanism for team AI. This mechanism, called TEAM, is discussed next.

3 Tactics Evolutionary Adaptability Mechanism (TEAM)

The Tactics Evolutionary Adaptability Mechanism (TEAM) is an evolutionary inspired adaptation mechanism that imposes adaptive team AI on opponents in games. In this chapter, the concept of TEAM is first laid out. Then, four features are discussed which distinguish TEAM from typical evolutionary approaches. These four features are: (1) a centralised agent control mechanism evolution, (2) a mutualistic-symbiotic evolution, (3) a delayed FSM-based evaluation, and (4) an evolution with history fall-back. A popular team-oriented game, the Quake III Capture The Flag (CTF) team-based game mode [7], is used for illustrative purposes.

3.1 Concept

TEAM is designed to be a generic adaptive mechanism for team-oriented games in which the game state can be represented as a finite state machine (FSM). An instance of TEAM is created for each state of the FSM. Each instance, in fact, is an evolutionary algorithm which learns state-specific behaviour for *a team as a whole*. In our experiment, the evolutionary algorithm was designed to learn optimal parameter values for each state's team control mechanism.

Cooperatively, all instances of TEAM learn successful team-oriented behaviour for *all states*. The concept of TEAM is illustrated in figure 2.

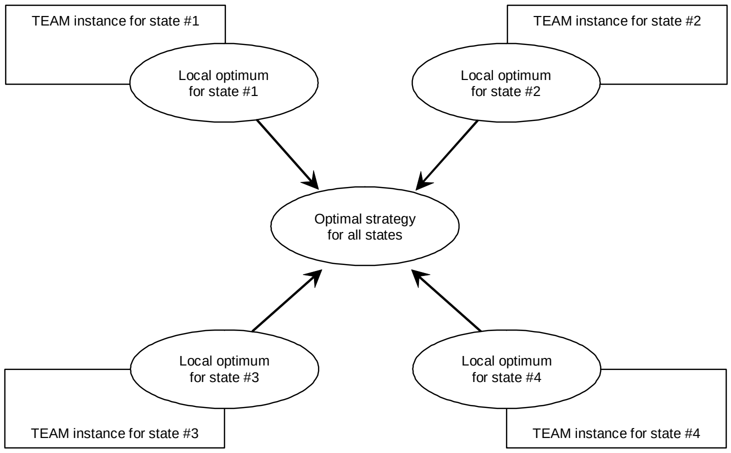

Fig. 2. Conceptually, TEAM learns adaptive behaviour for *a team as a whole* (rather than learning adaptive behaviour for each individual). Instances of TEAM cooperatively learn team-oriented behaviour, which is defined as the combination of the local optima for the states (in this example there are four states).

3.2 Centralised agent control mechanism evolution

TEAM evolves the agent team control mechanism in a centralised fashion. The choice for the centralised approach is motivated by the desire to evolve the behaviour for a team as a whole. The performance of a team's behaviour is assessed by performing a high-level evaluation of the whole team.

3.3 Mutualistic-symbiotic evolution

The evolutionary approach of TEAM is inspired by mutualistic-symbiotic animal communities [8]. In such communities, individuals postpone short-term individual goals in favour of long-term community goals.

The focus of TEAM's evolutionary mechanism lies on learning team-oriented behaviour by the cooperation of multiple evolutionary algorithms. For each state of the FSM controlling the team-oriented behaviour, a separate evolutionary algorithm is used. Each of these evolutionary algorithms learns relatively uncomplicated team-oriented behaviour for the specific state only. Yet, the behaviour is learned in consideration of the long-term effects on the other evolutionary algorithms. Subsequently, relatively complex team-oriented behaviour emerges in a computationally fast fashion.

3.4 Delayed FSM-based evaluation

We defined a delayed FSM-based evaluation function that postpones the determination of the fitness value of a genome until a certain number of state transitions (the so-called "depth") after employing the genome have been processed. The delayed FSM-based evaluation function consists of two components: (1) a scaled fitness function, and (2) a delayed fitness function.

First, the scaled fitness function is defined as:

$$Scaled\_fitness = \begin{cases} 1.0 - \min\left(\frac{\sqrt{t}-\sqrt{\frac{t}{3}}}{10}, 1.0\right) & \{transition = +\} \\ \min\left(\frac{\sqrt{t}-\sqrt{\frac{t}{3}}}{10}, 1.0\right) & \{transition = -\} \end{cases} \tag{1}$$

where t denotes the time before a state transition occurs, and $transition$ is calculated by using annotations on the FSM which describes the states of a game. In figure 3, an example of annotations on the FSM of Quake III CTF is given. To realise the time-scaling, a damping square root is used, which has a substantial effect on short state transitions, and a clearing effect on long state transitions.

Second, the delayed fitness function for state transition M is defined as:

$$Delayed\_fitness_M = \sum_{i=0}^{n} \frac{1}{i+1} \left(Scaled\_fitness_{M+i}\right) \tag{2}$$

where i is the depth, n is a positive integer, and $Scaled\_fitness$ is the scaled fitness value of a genome. The delayed-reward is used to consider the long-term effects of genomes, because positive behaviour is only desirable if the team can retain or improve on the behaviour. In our experiment a two-deep delayed-reward is used ($n = 2$).

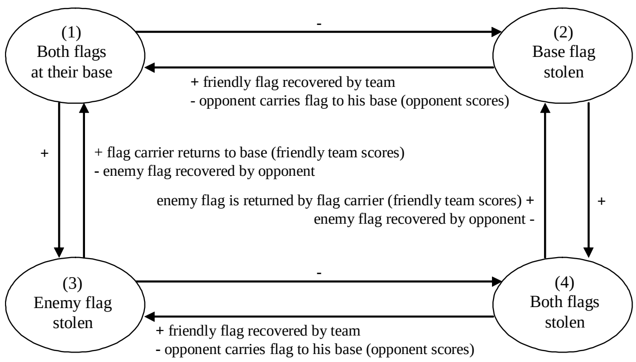

Fig. 3. Annotated finite state machine of Quake III CTF. Desirable state transitions are denoted with "+", whereas undesirable state transitions are denoted with "−".

3.5 Evolution with history fall-back

The game-environment of team-oriented games is typically accompanied by a large amount of randomness. The randomness poses a problem for most adaptive mechanisms since one cannot be sure that a successful course of the evolution is the direct result of the genetic information in the population, or, of lucky circumstances. Consequentially, the evolutionary mechanism of TEAM is capable of reverting to an earlier state, if this is required. We implemented history fall-back with a fitness recalculation mechanism, which filters out unsuccessful genomes in due time.

4 Experiment: TEAM vs. Quake III team AI

We tested the TEAM adaptation mechanism in a Quake III CTF game, where an adaptive team is pitted against a non-adaptive team. The adaptive team is controlled by TEAM. The non-adaptive team is controlled by the Quake III team AI, which offers non-static and fine-tuned, but non-adaptive, behaviour. The Quake III team AI changes its behaviour in such a way, that the adaptive team has to be able to deal with significant behavioural changes of the opponent. Both teams consist of four agents with identical individual agent AI, means of communication and team organisation, and only differ in the control mechanism employed (adaptive or non-adaptive).

4.1 Evaluation of an experimental run

An experimental run consists of two teams playing Quake III CTF until the game is interrupted by the experimenter. To quantify the performance of TEAM, two properties of an experimental run are used: the absolute turning point and the relative turning point.

We define the absolute turning point as the time at which the adaptive team obtains a win-loss ratio of a least 15 wins against 5 losses in a sliding window of 20. When the ratio is reached, the probability of the adaptive team outperforming the non-adaptive team is $> 98\%$ [9].

We define the relative turning point, which quantifies the noticeable effect of successful adaptive behaviour, as the last time at which the adaptive team has a zero lead with respect to the non-adaptive team, with the additional requirement that from this moment on the adaptive team does not lose its lead for the rest of the experimental run.

4.2 Results

In table 1 an overview of the experimental results is given. The average absolute turning point acquired is 108, and the average relative turning point is 71. TEAM merely requires several dozens of trials to evolve excellent behaviour, which is a good result especially considering that evolutionary algorithms typically require several thousands of trials to achieve successful results.

	Absolute turning point	**Relative turning point**
Avg	108	71
StDev	62	45
StError	19	14
Median	99	50
Minimum	38	20
Maximum	263	158
Avg-StError	89	57
Avg+StError	127	85

Table 1. Summary of experimental results.

To illustrate the course of a typical experimental run, we plotted the absolute and relative performance in figure 4. As shown in the top graph of figure 4, initially the adaptive team obtains approximately 10 wins against 10 losses; this is considered to be neutral performance. At the absolute turning point (point 99) a dramatic increase of the performance of the adaptive team is observed. In the bottom graph of figure 4, we observe that, initially, the adaptive team attains a lead of approximately zero. At the relative turning point (point 50) the lead of the adaptive team dramatically increases.

4.3 Evaluation of the results

The experimental results imply that TEAM is able to successfully counter non-static opponent behaviour, as it defeated the non-static Quake III team AI. In the next section we will argue that this is the result of TEAM discovering and applying unforeseen dominant tactics.

Moreover, table 1 shows that the average relative turning point is much below the average absolute turning point. From this observation we may conclude that before we can reliably determine if the absolute turning point is reached, the opponent already notices the dominant effect of TEAM.

Considering that in all runs we were able to determine relatively low absolute- and relative turning points, implying that TEAM learned to significantly out-perform the opponent (in this case the Quake III team AI), we may draw the conclusion that TEAM is capable of successfully adapting to significant changes in the opponent behaviour. Figure 4 (below) reveals that TEAM learned to out-perform the opponent without any significant degradation in the adaptive team's performance.

5 Discussion

Our experimental results show that TEAM is a successful adaptive mechanism for team-oriented behaviour. In sub-section 5.1 we will discuss to what extent

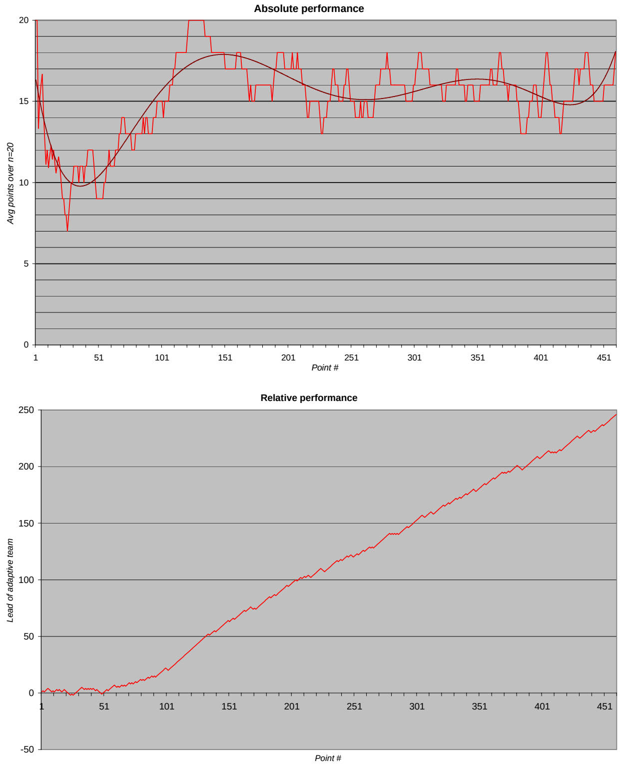

Fig. 4. Illustration of typical experimental results obtained with TEAM. The top graph shows the points scored by the adaptive team over a sliding window of 20 as a function of the amount of scored points. The bottom graph shows the lead of the adaptive team over the non-adaptive team as a function of the amount of scored points. The bottom graph reveals that the adaptive team outperforms the non-adaptive team without any significant degradation in its performance.

TEAM meets the requirements necessary to allow it to be implemented in actual games. The behaviour learned by TEAM is discussed in sub-section 5.2.

5.1 Qualitative evaluation of TEAM

TEAM is an online adaptive mechanism. For online adaptation to work in practice, we denote four requirements for qualitatively acceptable performance. It must be (1) computationally fast, (2) robust with respect to randomness inherent in the environment, (3) efficient with respect to the number of adaptation trials, and (4) effective with respect to the intermediate AI generated during the adaptation phase [10].

TEAM is computationally fast and is able to cope with a large amount of randomness inherent in the environment [11]. As argued in section 4.3, TEAM is efficient with respect to the limited number of trials required for an opponent to notice the effects of dominant adaptive behaviour. The effectiveness of TEAM is expressed by the fact that it outperforms non-adaptive AI without any significant degradation in performance.

We may therefore conclude that TEAM is computationally fast, robust, efficient and effective, and can be applied in practice.

5.2 Learned behaviour

Analysing the behaviour of TEAM, we observed that the population does not converge to merely one set of dominant tactics. TEAM is continuously adapting to the environment in order to remain dominant throughout the game. In our experiment, the adaptive team has learned risky, yet successful, tactics. These tactics can be best described as "rush" tactics, which are often applied in the real-time strategy game genre. Rush tactics aim at quickly obtaining offensive field supremacy. If rush tactics are successful, the opponent can seldom recover from the momentum of the offensive team.

The original Quake III team AI uses only moderate tactics in all states. Therefore, it is not able to counter *any* field supremacy. This exemplifies the inadequacy of non-adaptive AI. Despite the fact that the original Quake III team AI is fine-tuned to be suitable for typical situations, it cannot adapt to superior player tactics, whereas TEAM can.

6 Conclusions and Future Work

We proposed the Tactics Evolutionary Adaptability Mechanism (TEAM) as a novel mechanism to impose adaptive behaviour on opponents in team-oriented commercial computer games. TEAM is based on evolutionary algorithms, but it possesses four features which distinguish the adaptive mechanism from typical evolutionary approaches. These four features are: (1) a centralised agent control mechanism evolution, (2) a mutualistic-symbiotic evolution, (3) a delayed FSM-based evaluation, and (4) an evolution with history fall-back. From the experimental results of our Quake III CTF experiment, we drew the conclusion that TEAM is capable of successfully adapting to changes in the opponent behaviour. TEAM adapted the team behaviour in such a way that dominant tactics were discovered. These dominant tactics outperformed the Quake III CTF tactics, which exhibit more cautious behaviour. Therefore, we may conclude that by creating a team-oriented mechanism capable of unsupervised and intelligent adaptation to the environment, we succeeded in creating challenging adaptive team AI.

Our future work aims at designing an adaptive mechanism capable of deciding which behavioural adaptations are required for endowing opponent AI with *entertaining* behaviour, and ultimately, designing a technique capable of learning meta-models for autonomously determining how to best accomplish an adaptation.

References

[1] Fenley, R.: Baby Steps to our Future. (2002) http://www.hal-pc.org/journal/02novdec/column/babystep/babystep.html

[2] FIPA: The foundation for intelligent physical agents. (2001) http://www.fipa.org

[3] Tozour, P.: The Perils of AI Scripting. AI Game Programming Wisdom (ed. Rabin, S.). Charles River Media. ISBN: 1-58450-077-8 (2002) 541–547

[4] Mitchell, T.M.: Machine Learning. McGraw-Hill Science/Engineering/Math. ISBN: 0-070-42807-7 (1997) 368–390

[5] Eiben, A.E.: Smith, J.E.: Evolutionary Computing Tutorial Part 1: What is an Evolutionary Algorithm? (2003) http://www.evonet.polytechnique.fr/CIRCUS2/valise/bcraenen/intro_tut.pdf

[6] van der Sterren, W.: Squad Tactics - Team AI and Emergent Maneuvers. AI Game Programming Wisdom (ed. Rabin, S.). Charles River Media. ISBN: 1-58450-077-8 (2002) 233–246

[7] van Waveren, J.P.M., Rothkrantz, L.J.M.: Artificial Player for Quake III Arena. 2nd International Conference on Intelligent Games and Simulation GAME-ON 2001 (eds. Quasim Mehdi, Norman Gough and David Al-Dabass). SCS Europe Bvba. (2001) 48–55

[8] Raven, P.H., Johnson, G.B.: Biology. McGraw-Hill Science/Engineering/Math. ISBN: 0-073-03120-8 (2001)

[9] Cohen, P.R.: Empirical Methods for Artificial Intelligence. ISBN: 0-262-03225-2 (HC) (1995)

[10] Spronck, P.H.M., Sprinkhuizen-Kuyper, I.G., Postma, E.O.: Online Adaptation of Computer Game Opponent AI. Universiteit Maastricht, The Netherlands (2003)

[11] Bakkes, S.C.J.: Learning to Play as a Team. Master's thesis. Universiteit Maastricht, The Netherlands (2003) 41–45