

Honest vs Cheating Bots in PATROL-based Real-Time Strategy MMOGs

Stefano Sebastio, Michele Amoretti, Jose Raul Murga,
Marco Picone, and Stefano Cagnoni

Abstract The increasing success of massively multi-player online games (MMOGs) is due to the fact that they allow players to explore huge virtual worlds and to interact in many different ways, either cooperating or competing. To support the implementation of ultra-scalable real-time strategy MMOGs, we are developing a middleware, called PATROL, that is based on a structured peer-to-peer overlay scheme. Among other features, PATROL provides AI-based modules to detect cheating attempts, that the decentralized communication infrastructure may favor. In this paper we illustrate how we implemented honest and cheating autonomous players (bots). In particular, we show how honest bots can detect cheating bots in real-time, thanks to strategies based on neural networks.

Stefano Sebastio
IMT Institute for Advanced Studies Lucca, Italy, e-mail: stefano.sebastio@imtlucca.it

Michele Amoretti
Centro Interdipartimentale SITEIA.PARMA,
Università degli Studi di Parma, Italy, e-mail: michele.amoretti@unipr.it

Jose Raul Murga
Dip. di Ingegneria dell'Informazione,
Università degli Studi di Parma, Italy, e-mail: josemurgav@gmail.com

Marco Picone
Dip. di Ingegneria dell'Informazione,
Università degli Studi di Parma, Italy, e-mail: picone@ce.unipr.it

Stefano Cagnoni
Dip. di Ingegneria dell'Informazione,
Università degli Studi di Parma, Italy, e-mail: stefano.cagnoni@unipr.it

1 Introduction

Most research on multi-player online games (MMOGs) focuses on scalability and high speed, but other issues such as the chance of cheating have an equally large practical impact on game success. There are significant technical barriers to achieving all these properties at the same time, and few existing games do so. Our open source middleware for the development of real-time strategy (RTS) MMOGs, called PATROL (<http://code.google.com/p/patrol/>), integrates several modules, each of which is highly specialized in one aspect of the game.

In our previous work [1] we focused on the module that enables peer-to-peer connectivity and communication, and on the module for detecting cheating behaviors. The former allows one to implement ultra-scalable RTS MMOGs, where each player's software installation is a node of a fully distributed structured overlay network scheme, which guarantees efficient data sharing, as well as fair and balanced workload distribution among participants. The latter, based on MLP neural networks, allows each node to detect other players' cheating behaviors.

In this paper we present the Rule Engine, the PATROL module that allows one to enforce both general and specific rules into the nodes, and matches game events with existing rules. The Rule Engine allowed us to implement autonomous playing agents, called *bots*, that are able to play one against another, in an RTS game where participants place and move units and structures (generally speaking, *resources*) to secure areas of the virtual world and/or destroy the assets of their opponents. We show how honest bots are able to detect cheating ones in real-time, thanks to the cheating detection module.

The paper is organized as follows. In section 2 we summarize some recent research work in the context of peer-to-peer RTS MMOGs. In section 3 we describe the PATROL architecture, with details on the rule engine and on the cheating detection strategy. In section 4 we present the example of RTS MMOG we have implemented, based on the proposed architecture, where autonomous bots play one against another. In section 5 we illustrate experimental results, focusing on the performance of the module for intelligent cheating detection, presenting and discussing many experimental results. Finally, in section 6, we conclude the paper by specifying plans for extending our work further.

2 Related Work

MMOG needs a messaging infrastructure for game actions and player communication. To this purpose, possible paradigms are Client-Server (CS), Peer-to-Peer (P2P), Client-Multi-Server (CMS), or Peer-to-Peer with Central Arbiter (PP-CA) [2]. Each solution has pros and cons, with respect to robustness, efficiency, scalability and security. In particular, when the architecture of the game is decentralized (*e.g.* P2P), with a large number of players, then facing malicious behaviors is particularly challenging.

In [2], the authors propose a Mirrored-Arbiter (MA) architecture that combines the features of CMS and PP-CA. This architecture provides all the benefits of PP-CA, but also solves the main problems in PP-CA by using interest management techniques and multicast. Clients are divided into groups, each group being handled by an arbiter that maintains a global state of the game region and takes care of the consistency issue. When the arbiter receives an update from a client which conflicts with its game region state, it ignores the update and sends the correct region state to all clients in the group. The authors implemented a multiplayer game called "TankWar" to validate the design of the proposed MA architecture. In our opinion, such a scheme is complex in the decision of the arbiters and their group assignments, and does not guarantee high scalability and security. Indeed, an arbiter may be a cheating node itself, which compromises the game for a large number of nodes.

In [3], the authors present a Peer-to-Peer (P2P) MMOG design framework, Mediator, based on a super-peer network with multiple super-peer (Mediator) roles. In this framework, the functionalities of a traditional game server are distributed, capitalizing on the potential of P2P networks, and enabling the MMOG to scale better with respect to both communication and computation. Mediator integrates four elements: a reward scheme, distributed resource discovery, load management, and super-peer selection. The reward scheme differentiates a peer's contribution from its reputation, and pursues symmetrical reciprocity as well as discouraging misdemeanors. The authors suggest to adopt the EigenTrust reputation management algorithm [5] and the DCRC anti-free-riding algorithm [6] as possible implementations for the reward scheme. Unfortunately, such schemes are complex and bandwidth-consuming.

The main aspect on which our paper focuses is the game engine. In our opinion, games are made of *nouns* (*i.e.* elements of the game, and variables related to them), *verbs* (the actions that players and player stand-ins can enact), and *rules* (limiting the scope of the nouns and creating relationships and interactions between them; limiting also which verbs can be enacted, when and in which context). Current programming paradigms do not provide an appropriate language for expressing these structures. Object-Oriented Programming (OOP) is very good at representing different types of objects ("nouns") and the relationships they may have between each other, with minimal duplicated code or wasted work. Unfortunately, OOP dictates that each class must encapsulate methods, *i.e.* actions that are strictly related to the class itself. Thus, OOP is not suitable for expressing verbs and rules as entities separated from nouns. In general, imperative languages (that are mostly used in game programming) are not good for clearly expressing verbs and rules. Instead, declarative languages, like Prolog, are much more suitable.

Currently, one of the best known rule engines is Drools Expert [7], that uses the rule-based approach to implement an Expert System and is more correctly classified as a Production Rule System. A Production Rule System is Turing complete, with a focus on knowledge representation to express propositional and first order logic in a concise, non-ambiguous and declarative manner. The core of a Production Rule System is an Inference Engine that is able to scale to a large number of rules and

facts. The Inference Engine matches facts and data against Production Rules — also called Productions or just Rules — to infer conclusions which result in actions. A Production Rule is a two-part structure that uses First Order Logic for reasoning over knowledge representation. There are a number of algorithms used for Pattern Matching by Inference Engines including Linear, Rete, Treat. Drools implements and extends the Rete algorithm and is a sound product, but it is quite large and cannot be installed on portable devices. For this reason, we decided to adopt tuProlog [8], as discussed in the next section.

3 PATROL middleware

In order to increase security, the game infrastructure should properly manage the interaction events among nodes. In RTS games, the most frequent events are those for: i) moving resources, ii) receiving updates about the virtual world, and iii) submitting the attacks. Our PATROL middleware manages these events through protocols that are appropriate for maintaining an adequate level of efficiency and security.

PATROL provides the following modules (illustrated in Fig. 1):

- Rule Engine
- Cheating Detector
- Overlay Manager
- GUI/GamePeer Connector

Since the Overlay Manager and the Cheating Detector have been already described in details in our previous work [1], here we just recall them shortly and we devote more space to the Rule Engine. The GUI/GamePeer Connector decouples the (game-specific) GUI from the GamePeer, which integrates the three previously listed general-purpose modules. For lack of space, we omit its description, but we emphasize that GUI decoupling also allows one to implement games for mobile devices, where only the visualization may be running locally, while most computation processes may be executed remotely.

3.1 *Overlay Manager*

PATROL's Overlay Manager adopts the Chord P2P overlay scheme [9] to support fair and robust information sharing among available players. Chord is a highly structured P2P architecture where all peers are assigned the same role and amount of work. It is based on the DHT approach for an efficient allocation and recovery of resources. The overlay network in PATROL also supports a distributed algorithm for cheater detection, based on feedbacks among peers and AI tools such as neural networks. This approach allows one to dynamically recognize malicious behaviors,

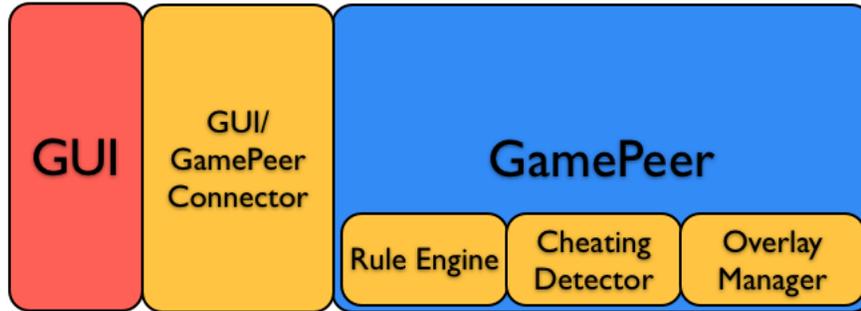


Fig. 1 A PATROL-based gaming node: PATROL modules are those in light color.

collectively performed by peers without the need of specific and centralized control components.

Chord [9] is probably the best known peer-to-peer protocol based on the Structured Model (SM), which uses DHTs as infrastructures for building large scale applications. Data are divided into *blocks*, each identified by a unique *key* (a hash of the block's name) and described by a *value* (typically a pointer to the block's owner). Each peer is assigned a random ID in the same space of data block keys, and is responsible for storing (key,value) pairs for a limited subset of the entire key space.

According to Chord's lookup algorithm, each node n maintains a routing table with up to m entries, called the *finger table*. The i^{th} entry in the table at node n contains the identity of the first node s that follows n by at least 2^{i-1} on the identifier circle; i.e. $s = \text{successor}(n + 2^{i-1})$, where $1 \leq i \leq m$ and all the arithmetic is module 2^m . We call node s the i^{th} *finger* of node n , and denote it by $n.\text{finger}[i]$. A finger table entry includes both the Chord identifier and the IP address (and port number) of the relevant node. Fig. 2 illustrates the scalable lookup algorithm based on finger tables. In general, if node n searches for a key whose ID falls between n and the successor of n , node n finds the key in the successor of n ; otherwise, n searches its finger table for the node n' whose ID most immediately precedes the desired key, and then the basic algorithm is executed starting from n' . It is demonstrated that the number of nodes that must be contacted to find a successor in an N -node network is $O(\log N)$ [9] in the majority of cases.

PATROL distributes uniformly among the peers the responsibility to maintain knowledge about resources (i.e. items, war units, structures, etc.), using the DHT to share information about whom is responsible for what (each peer is responsible for a subset of the key space). In a game, each existing resource has a position in the virtual world. Such a position is hashed, and the resulting key is assigned to the peer whose key subset includes the resource key. It is very unlikely that two resources that are close in the virtual world have keys that are also close in the key space (and vice versa). Moreover, Chord provides data replication, in order to improve robustness against unexpected node departures.

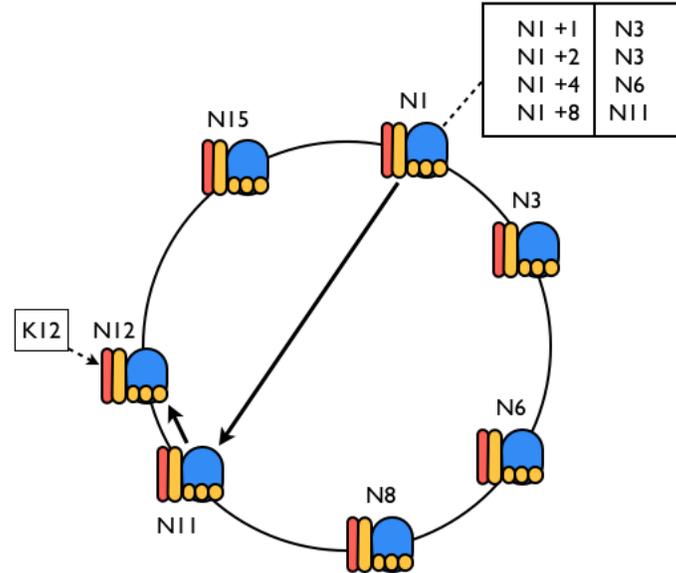


Fig. 2 The finger table entries for node $N1$ and the path taken by a query from $N1$, searching for key $K12$ using the scalable lookup algorithm.

As a consequence, the control over a region of the game space is distributed among several peers, thus limiting the damages that a hacked peer can do. Moreover, our approach is much more robust than existing decentralized solutions, because the departure of a node does not affect the games too much, thanks to the proactive data replication dictated by the Chord protocol.

3.2 Rule Engine

In general, a rule engine is a software system that, depending on the context, decides which rules to apply, and computes the result of their application, that may be a new knowledge item, or an action to perform. Usually, a rule engine includes the following components:

- a *rule base*, containing production rules whose structure is **WHEN** ⟨conditions⟩ **THEN** ⟨actions⟩;
- a *knowledge base*, also known as *work area*, that contains known facts;
- an *inference engine*, for processing rules.

Rules operate on facts of the knowledge base, that is dynamic as it can change over time, with new facts being added, and old facts being removed. Conditions of production rules are evaluated against facts. If a condition is true, the resulting

action is the insertion of a new fact in the knowledge base, and possibly an action on the environment (*e.g.*, an action within the game).

PATROL's Rule Engine (from now on, RE) can be used for implementing several RTS MMOGs: it is sufficient to set the appropriate rule base and knowledge base. Rules and facts must be written in Prolog, chosen because of its intuitiveness. RE's inference engine is based on tuProlog, a Java-based lightweight Prolog reasoner for Internet applications and infrastructures [8]. tuProlog provides a straightforward API to embed Prolog programs within Java code, or to read existing Prolog expressions from a file or from a database. Once one or more Prolog *theories* (*i.e.* ensembles of rule base and initial knowledge base) have been acquired, it is possible to use them to evaluate facts and derive new facts.

The RE can be used to decide which actions are allowed to the player, depending on his/her state and on the state of the game. Moreover, the RE can be used to implement bots, *i.e.* autonomous playing agents that allow, for example, to test game strategies before entering a game against other real players. A bot must be able to make decisions in all typical RTS situations, such as: resource accumulation, resource purchasing or building, resource improvement, displacement of mobile resources, attack against an enemy's resources, defense from an enemy's attack, goal checking.

The RE includes a `PrologEngine` class that provides methods for setting and managing a theory, and for solving queries. Such a class can be specialized (by means of inheritance) into different classes, each referring to an aspect of the game. Such specialized classes can be reused with different theories, and within different RTS MMOGs.

3.2.1 Game Events

The system uses a bootstrap server to support peers in joining the network (which includes authentication, as well as Chord initialization) and configuring themselves for entering a game. In this way the bootstrap server has control over the accounts of the players and consequently provides a basic level of security.

Information about the virtual world may not be granted indiscriminately to any peer. Each peer has its own resources, which are placed in different positions of the virtual world, and has the right to receive information that refers to areas that are within the field of view of such resources, according to the rules of the game. Periodically, each peer needs to update its view on the virtual world. To do so, it sends specific requests to peers that are responsible for the positions that are visible. Before responding to such a request, peer N_j , that we suppose to be responsible for position (x, y, z) , checks its cache for updated information, and sends a request to verify the credentials for another peer N_k . If everything is ok, it finally sends the response message.

To perform any action that involves a change of game state, players must submit a request to the responsible of the location that is affected by the action (like in figure 3). For example, suppose that a peer N_k can select a resource to be displaced in the

virtual world; to perform the action "displacing resource to position (x,y,z) " the peer must submit the request to the node responsible for the key resulting from the hash of that position, *i.e.* $h(x,y,z)$. The peer j that must become the new responsible for the displaced resource of peer N_k searches for the manager of the resource's current position (declared by peer N_k). Such peer is discovered by means of the hash of the current position. Thus the old manager checks in its cache whether it has the information on peer N_k and whether this information corresponds to what was declared to N_j . If the check is successful, peer N_j can decide, according to the game rules and considering the time elapsed between the changes of game state following the transition between the two positions, if it can accept the move and execute it, becoming the new responsible for the resource. If the position declared by N_k is not true, the request for resource displacement submitted by peer N_k is ignored and the state of the game remains the same.

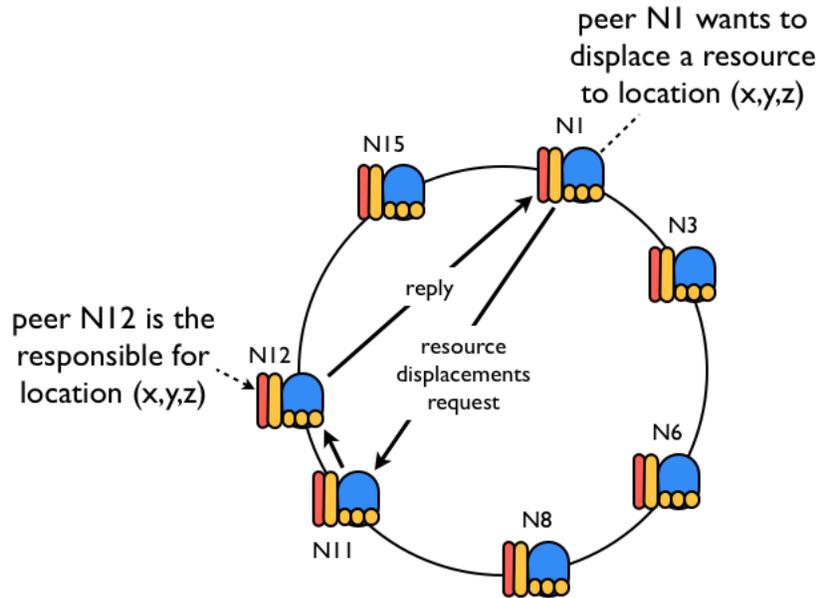


Fig. 3 Submission of an action request to the responsible of the location that is affected by the action.

While troops can be moved asynchronously by each player, attacks can be either asynchronous or synchronous (*i.e.* with a turn-based approach). In case of synchronous attacks, knowing the decisions of other players before submitting his own move may be a considerable advantage for a player. But, of course, this would be unfair.

To avoid cheating, PATROL uses request hashing, a mechanism that is widely adopted in other P2P architectures and derives from distributed security systems. Players who submit their decisions have to send a hash of the message describing

the attack concatenated with a *nonce*. The nonce is used to prevent a cheater from storing in a table all matches between hash values and attack decisions, revealing the decisions of honest players. The nonce is a use-once random value, chosen by the first player that submits a decision. The last player that submits its decision may send it manifestly. At this point, all players that have previously sent the hash must send their nonce and attack decision manifestly. Thus, other players can re-calculate the hash and verify that it corresponds to what was previously declared. The properties of hash functions guarantee that it is almost impossible for two different attacks to have the same hash, and therefore for a player to submit a different attack, with respect to the hashed one.

3.3 Intelligent Cheater Detection

PATROL provides a good level of security for the overall state of the game. However, the DHT does not prevent the game from offering cheaters (provided with hacked clients) the possibility to alter the information for which they are responsible. In an RTS game, a modified client that saves a history of recent attacks and their outcomes may estimate the current level of resources available to other players and take advantage over them.

Using artificial intelligence techniques, a PATROL-based peer can detect anomalies in the behavior of other peers, compared to typical behavioral profiles, by means of temporal analysis of interaction events. Moreover, using the power of direct communication typical of P2P approaches, a peer may ask other peers their opinion about a given peer in order to improve the evaluation process.

Peer x calculates the probability $P\{y|x\}$ that peer y is cheating. Then x sends a request to peers that have interacted with y , in order to match their probabilities and understand whether y is considered to be a cheater: $P\{y|i\} \forall i \neq x, y$. If the global probability exceeds a certain threshold, there is the option to contact other peers and the bootstrap server to promote a collective motion against the cheating player in order to ban him from future games. If all peers agree with the "Ban Proposal", peer y is gracefully disconnected from the Chord ring.

The Cheating Detector module (Fig. 1) analyzes all action events coming from an opponent. The opinions of the other players are requested only at the end of the local evaluation process, if the peer estimates a high probability of cheating. Of course, the peer must be careful since other peers may provide false reports related to their interactions with a given peer.

There are different strategies for a peer to learn from a sequence of events: sequence recognition, sequence playback, and temporal association. Among others, we focused on *Multi Layer Perceptrons (MLPs)*, that we implemented by means of the Weka library [4]. We analyzed their features in details in our previous work [1].

MLPs are the traditional multilayer neural networks trained by the backpropagation algorithm. Weights are updated using the online algorithm, *i.e.* re-arranging the weights after each epoch, *i.e.* a learning iteration during which all examples

are processed by the network. Neural networks have been used in similar contexts, *e.g.* for the detection of cheating players in first-person shooters [10], but also with slightly different objectives, like the detection of bots that play in place of human participants [11, 12].

4 Game Bots

We have extended the `PrologEngine` class of the RE into specialized classes, to implement a RTS MMOG with space setting. In such a RTS MMOG, players have to find and conquer all the planets that are in the game space.

Players are provided with a mine resource that allows them to make money for buying two types of resources: defense and attack. The resources for attack (starships) are used to explore the virtual world, to the purpose of finding the planets and to tackle the starships of other players. Every resource has an associated velocity and a field of view. The resources for defense are used to protect the owned planets from incoming attacks of other players' starships.

Thus, we have implemented the following modules (corresponding to Java classes): `ExtractionEngine` for managing the extraction of mineral resources of a planet, `BuyResourceEngine` for purchasing resources, `MovementEngine` for moving mobile resources (like starships), `VisibilityEngine` for deciding if a resource (*e.g.* a planet, or an enemy's starship) is visible to the player, `GameEvolutionEngine` for deciding the next operation depending on current state (own state, and game state), `GameEngine` for checking if intermediate or final goals have been met.

For testing purposes, we have implemented *game bots*, *i.e.* virtual players that automatically handle game tasks — extracting mineral resources, purchasing resources, moving mobile resources, etc. A game bot is a weak AI agent which for each instance of the program, may play against other bots and/or human players at the same time, either over the Internet, on a LAN or in a local session [13]. Advanced bots feature machine learning for dynamic learning of patterns of the opponent as well as dynamic learning of previously unknown maps. Using bots to control characters that are supposed to be controlled by human players is incidentally against the rules of all of the current main Massively Multiplayer Online Role-Playing Games (MMORPGs), such as Ultima Online and World of Warcraft. However, the virtual worlds of all MMOGs are characterized by a number of *non player characters (NPC)* with different objectives — either collaborative or competitive with regard to players. For this reason, they can be considered as artificial life systems.

Based on the previously listed engines, game bots pass through three different phases:

1. resource accumulation
2. space exploration
3. planet conquest

These phases are repeated in an infinite loop. The time periods each bot spends in such phases are random variables.

We have defined two different types of bot profiles: honest and cheater. The latter reproduces the behavior of a hacked client. It owns a mine which is five times as powerful as the others and more initial money. Moreover, it has a halved cycle decision period (*i.e.* it can take more decisions in the same time).

In Fig. 4, we report the distribution of honest and cheating bots' actions during the recorded matches. On the horizontal axis we have the time at which the action is performed, considering 0 as the start time of the game, while on the vertical axis there is the value associated to the used resource. Here we can note that cheater bots, thanks to the speed hack, prefer to first explore the space and to perform their actions later than the honest bots. Moreover, the value associated to their resources is higher since the more money is available, the more they can spend for buying higher-valued resources.

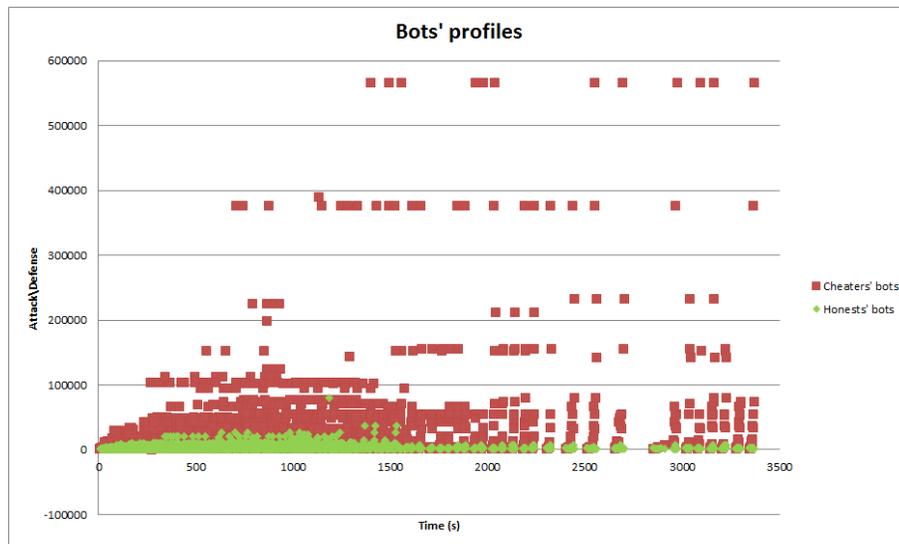


Fig. 4 Time distribution of bots' actions, considering all the games that have been played to collect a significant number of actions.

5 Experimental evaluation

We have collected 2800 player profiles, each one consisting of a sequence of three actions performed by the player. The choice of the three moves for the profile duration appears as a good trade-off between quickness and accuracy of evaluation of the opponent's behavior. The overall dataset has been split into a training set of

1300 profiles, a validation set of 100 profiles and a test set of 1400 profiles. For all these sets, 50% are profiles related to honest players, and 50% are profiles related to cheating players.

We have defined different configurations for the MLP neural networks, that we have compared these by evaluating their *Root Mean Square Error (RMSE)*. All the MLP neural networks we have considered have three layers, with x, y, z neurons each. In particular, we have used the following configurations: 6, 6, 1; 6, 5, 1; 6, 4, 1; 6, 3, 1; 6, 2, 1; 6, 1, 1. Such networks have been compared by measuring their RMSE with respect to the validation set. Those that have shown the best performances have been selected and tested using the test set. For all six configurations we have considered different values for the seed, the momentum, the learning rate and the epoch number. In details, we have used 5 seeds over which we have computed the average, variance and standard deviation of the RMSE. We have also computed the 95% confidence interval using Student's T-test. The momentum and the learning rate have been varied with step equals to 0.1 in the range $[0, 1]$, and for the epoch number we considered the following values: $\{5000, 10000, 20000, 30000, 40000, 50000, 60000\}$.

Fig. 5 and fig. 6 illustrate the RMSE values of the best MLP neural networks we have found during the experiments, considering the validation set and the test set. Fig. 5 refers to experiments where the networks are requested to provide a "cheating indicator" $\in [0, 1]$. Instead, the results in fig. 6 are related to experiments where the output of the network is a boolean (honest or cheating). Considering these, we observe that the most performant MPL networks are those with smaller RMSE values and standard deviations on the test set, *i.e.*:

- MLP 6, 6, 1(0.1; 0.0; 30000) – 42: 0.35188 ± 0.006782
- MLP 6, 3, 1(0.2; 0.0; 5000) – 21: 0.35176 ± 0.005714

where MPL $x, y, z(lr; m; e) - f$ means MLP with three layers of x, y, z neurons, respectively, learning rate lr , momentum m , epoch number e and f weights.

In general, MPL networks have yielded good performance on the test set, in terms of error percentage: detected cheating actions are between the 83% and 85% of the total number of actions during a match.

6 Conclusions

In this work we illustrated the most recent improvements of our PATROL framework for creating peer-to-peer online RTS games, characterized by a high degree of robustness, efficiency and effectiveness against cheating behaviors. In particular, we focused on the rule engine that allows us to enforce the rules of a game, and also to develop autonomous virtual players (bots). We have shown how cheating bots can be detected by means of a PATROL module that uses neural networks.

Tests have been encouraging, since cheating detection has a success rate of 85%. Moreover, it is possible to envisage the use of other means of temporal analysis

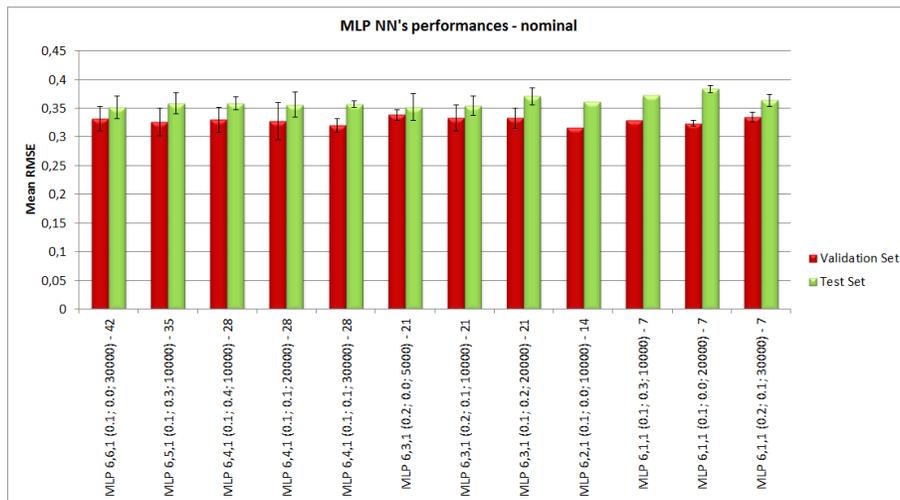


Fig. 5 The best RMSE values obtained with the MLP neural networks requested to provide a "cheating indicator" $\in [0, 1]$.

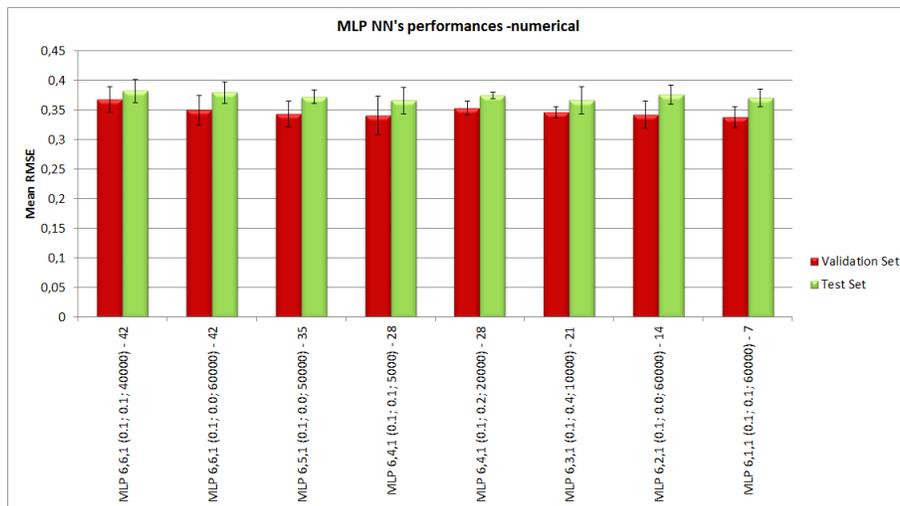


Fig. 6 The best RMSE values obtained with the MLP neural networks, where the output is a boolean (honest or cheating).

based on neural networks (such as *Real Time Recurrent Learning* and *Context Units* like Elman and Jordan nets) and on other techniques. It would also be possible to investigate the effects of adding a component capable of evaluating the *trust of peers* based on the past history of players.

Acknowledgements We would like to thank Dr. Alberto Lluch Lafuente of IMT Lucca for his precious comments that helped us to improve this work.

References

1. Picone M, Sebastio S, Cagnoni S, Amoretti M (2010) Peer-to-Peer Architecture for Real-Time Strategy MMOGs with Intelligent Cheater Detection. In: Proc. of Distributed Simulation & Online gaming (DISIO), co-located with 3rd ICST/ACM International Conference on Simulation Tools and Techniques (SIMUTools 2010), Torremolinos, Spain, March 2010.
2. Yang L, Sutinerk P (2007) Mirrored Arbiter Architecture – A Network Architecture for Large Scale Multiplayer Games. In: Summer Computer Simulation Conference (SCSC 2007), San Diego, California, USA.
3. Fan L, Taylor H, Trinder P. (2007) Mediator: A Design Framework for P2P MMOGs. In: Proc. of NetGames'07, Melbourne, Australia.
4. Hall M, Frank E, Holmes G, Pfahringer B, Reutemann P, Witten IH (2009) The WEKA Data Mining Software: An Update. SIGKDD Explorations, 11(1).
5. Kamvar SD, Schlosser MT, Garcia-Molina H (2003) The EigenTrust Algorithm for Reputation Management in P2P Networks. In: Proc. of the 12th Int'l Conf. World Wide Web, Budapest, Hungary.
6. Gupta M, Judge P, Ammar M (2003) A reputation system for peer-to-peer networks. In: Proc. of the 13th ACM NOSSDAV workshop, Monterey, CA, USA.
7. JBoss Community. Drools home page.
<http://www.jboss.org/drools>
8. Denti E, Omicini A, Ricci A (2005) Multi-paradigm Java-Prolog integration in tuProlog. Sci. Comput. Program., 57(2):217–250, 2005.
9. Stoica I, Morris R, Karger D, Frans Kaashoek M, Balakrishnan H (2001) Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications. In: Proc. of ACM SIGCOMM '01 Conference, San Diego, USA.
10. Galli L, Loiacono D, Cardamone L, Lanzi PL (2011) A Cheating Detection Framework for Unreal Tournament III: a Machine Learning Approach. In: Proc. of the IEEE Conference on Computational Intelligence and Games (CIG 2011), Seoul, South Korea.
11. Gianvecchio S, Wu Z, Xie M, Wang H (2009) Battle of Botcraft: Fighting Bots in Online Games with Human Observational Proofs. In: Proc. of CCS09, Chicago, Illinois, USA.
12. Prasetya K, Zheng W (2010) Artificial Neural Network for Bot Detection System in MMOGs. In: Proc. of the 9th Annual Workshop on Network and Systems Support for Games (NetGames '10), Taipei, Taiwan.
13. Kaminka GA, Veloso MM, Schaffer S, Sollitto C, Adobbati R, Marshall AN, Scholer A, Tejada S (2002) GameBots: A Flexible Test Bed for Multiagent Team Research. Communications of the ACM, 45(1):4345, 2002.