

The seal of the University of Bologna is a large, circular emblem in the background. It features a central figure, likely a saint or scholar, surrounded by various scenes and figures. The text "UNIVERSITAS BOLOGNENSIS" is written around the top inner edge, and "S. PETRUS UBIQUE PATER" is written around the bottom inner edge. At the very bottom, the word "SIGILLUM" is visible.

Fast and Fair Event Delivery in Large Scale Online Games over Heterogeneous Networks (Ph.D. thesis)

Claudio Enrico Palazzi

Technical Report UBLCS-2006-10

March 2006

Department of Computer Science
University of Bologna
Mura A. Zamboni 7
40127 Bologna (Italy)

The University of Bologna Department of Computer Science Research Technical Reports are available in PDF and gzipped PostScript formats via anonymous FTP from the area `ftp.cs.unibo.it:/pub/TR/UBLCS` or via WWW at URL `http://www.cs.unibo.it/`. Plain-text abstracts organized by year are available in the directory ABSTRACTS.

Recent Titles from the UBLCS Technical Report Series

- 2005-10 *Fault Tolerant Knowledge Level Communication in Open Asynchronous Multi-Agent Systems*, Dragoni, N., Gaspari, M., April 2005.
- 2005-11 *The AEDSS Application Ontology: Enhanced Automatic Assessment of EDSS in Multiple Sclerosis*, Gaspari, M., Saletti, N., Scandellari, C., Stecchi, S., April 2005.
- 2005-12 *How to cheat BitTorrent and why nobody does*, Hales, D., Patarin, S., May 2005.
- 2005-13 *Choose Your Tribe! - Evolution at the Next Level in a Peer-to-Peer network*, Hales, D., May 2005.
- 2005-14 *Knowledge-Based Jobs and the Boundaries of Firms: Agent-based simulation of Firms Learning and Workforce Skill Set Dynamics*, Mollona, E., Hales, D., June 2005.
- 2005-15 *Tag-Based Cooperation in Peer-to-Peer Networks with Newscast*, Marcozzi, A., Hales, D., Jesi, G., Arteconi, S., Babaoglu, O., June 2005.
- 2005-16 *Atomic Commit and Negotiation in Service Oriented Computing*, Bocchi, L., Ciancarini, P., Lucchi, R., June 2005.
- 2005-17 *Efficient and Robust Fully Distributed Power Method with an Application to Link Analysis*, Canright, G., Engo-Monsen, K., Jelasity, M., September 2005.
- 2005-18 *On Computing the Topological Entropy of One-sided Cellular Automata*, Di Lena, P., September 2005.
- 2005-19 *A model for imperfect XML data based on Dempster-Shafer's theory of evidence*, Magnani, M., Montesi, D., September 2005.
- 2005-20 *riends for Free: Self-Organizing Artificial Social Networks for Trust and Cooperation*, Hales, D., Arteconi, S., November 2005.
- 2005-21 *Greedy Cheating Liars and the Fools Who Believe Them*, Arteconi, S., Hales, D., December 2005.
- 2006-01 *Lambda-Types on the Lambda-Calculus with Abbreviations: a Certified Specification*, Guidi, F., January 2006.
- 2006-02 *On the Quality-Based Evaluation and Selection of Grid Services (Ph.D. Thesis)*, Andreozzi, S., March 2006.
- 2006-03 *Transactional Aspects in Coordination and Composition of Web Services (Ph.D. Thesis)*, Bocchi, L., March 2006.
- 2006-04 *Semantic Frameworks for Implicit Computational Complexity (Ph.D. Thesis)*, Dal Lago, U., March 2006.
- 2006-05 *Fault Tolerant Knowledge Level Inter-Agent Communication in Open Multi-Agent Systems (Ph.D. Thesis)*, Dragoni, N., March 2006.
- 2006-06 *Middleware Services for Dynamic Clustering of Application Servers (Ph.D. Thesis)*, Lodi, G., March 2006.
- 2006-07 *Meta Model Management for (Semi) Structured and Uncertain Models (Ph.D. Thesis)*, Magnani, M., March 2006.
- 2006-08 *Towards Abstractions for Web Services Composition (Ph.D. Thesis)*, Mazzara, M., March 2006.
- 2006-09 *Global Computing: an Analysis of Trust and Wireless Communications (Ph.D. Thesis)*, Mezzetti, N., March 2006.

Fast and Fair Event Delivery in Large Scale Online Games over Heterogeneous Networks (Ph.D. thesis)

Claudio Enrico Palazzi¹

Technical Report UBLCS-2006-10

March 2006

Abstract

The increasing need for people to be always connected, as well as the request for new amusement applications, poses several interesting issues in wireless networks. Nowadays, two main reasons are motivating several researchers around the world in striving to find effective solutions for these issues. First, revenues generated every year by electronic amusements follow a trend which is expected to increase further. Second, but certainly not less important, problems that emerge in developing innovative entertaining experiences are correlated with those belonging to various other “conventional” research fields in Computer Science.

In this context, online games for a large group of contemporary players are gaining attention as their participants are increasing in number. From a research point of view, they represent a very interesting and challenging topic especially in wireless environments. In this scenario, in fact, many issues arise such as authentication, session handling, interactivity, scalability and mobility, which are just a few and the more obvious of them.

Here, we propose a scalable architecture able to support fairness and interactivity in order to provide a cutting-edge online gaming experience to engaged players. In particular, we exploit a hybrid solution combining both the advantages of client-server and peer-to-peer paradigms. As part of our architecture, we have designed a novel synchronization mechanism among mirrored servers composing the online game platform. Our mechanism exploits the semantics of a game and follows a holistic approach to increase interactivity, fairness, and scalability, whilst maintaining full consistency of the game state. We provide extensive results that demonstrate our assertion.

However, even if our scheme is proficient in maintaining a high degree of responsiveness among the game servers, problems may arise at the edge of the considered topology. Concurrent traffic may generate queues that build up at the last (or first) link of a connection, thus delaying the game event delivery. This problem is worsened in case of players relying on wireless connectivity. The wireless medium, in fact, is naturally prone to be shared by several contemporary users interfering with each others' performance goals.

A typical case study is represented by a wireless home. Indeed, online entertainment is now possible at home by a plethora of ubiquitous services that can be provided based on wireless technologies, for instance, the common Wi-Fi IEEE802.11g network technology. Among the various home applications, entertainment is going to play a major role. As the availability of digital entertainment devices rapidly increases, the need for interconnecting them becomes even more urgent, as well as the necessity to extend the reach of entertainment to the wireless domain.

Applications run in this context may vary and, as we demonstrate in this thesis work, some of these may be particularly harmful toward real time traffic (i.e., online games, but also video-streaming and video-chats). In particular, we show how the very popular TCP-based FTP application for downloading files increases queuing delays to such an extent that responsiveness may be jeopardized.

¹ Department of Computer Science, University of Bologna, Mura A. Zamboni 7, 40127 Bologna, Italy.

To solve this problem, we propose and compare a set of possible solutions: i) appropriately setting MAC layer parameters, ii) employing TCP Vegas in place of the traditional TCP New Reno and, iii) limiting the advertised window for TCP flows. Extensive simulation results for a realistic in-home wireless scenario are provided.

Finally, we provide conclusion and future directions for our work. To this aim, it would be of particular interest to evaluate the IEEE 802.11e MAC protocol to investigate its capability to support real time applications and its impact on TCP-based traffic. Moreover, a particularly challenging case study is represented by the juxtaposition of our architecture with a vehicular networking scenario.

Contents

List of Figures	5
List of Tables	7
1 Introduction	8
1 Problem Statement and Thesis Contribution	<i>10</i>
2 Thesis Outline	<i>13</i>
2 Background	14
1 MMOG Issues	<i>14</i>
2 Fundamental Problems	<i>15</i>
3 Related Work	<i>17</i>
3.1 <i>Responsiveness</i>	<i>17</i>
3.2 <i>Consistency and Fairness</i>	<i>19</i>
3.3 <i>Cheating</i>	<i>20</i>
3.4 <i>Performance Evaluation over Wireless Links</i>	<i>21</i>
4 System Model	<i>21</i>
5 Architectures	<i>23</i>
3 Fast Synchronization Framework	26
1 Proposed Architecture	<i>26</i>
2 Obsolescence and Correlation: Maintaining Responsiveness and Consistency	<i>26</i>
3 Interactivity Restoring Mechanism	<i>28</i>
4 RED/RIO Techniques	<i>28</i>
5 Enhancing Interactivity with RED/RIO Techniques	<i>29</i>
6 Simulation Assessment	<i>33</i>
7 Current Results	<i>35</i>
7.1 <i>Obsolescence-Based Scheme vs Traditional One</i>	<i>35</i>
7.2 <i>ILA-RED vs ON-OFF: a Comparative Evaluation</i>	<i>36</i>
7.3 <i>ILA-RED: Sensitivity Analysis</i>	<i>38</i>
7.4 <i>ILA-RIO: Some Preliminary Results</i>	<i>39</i>
4 To Seek the Fairness by Way of the Interactivity	41
1 Exploiting Local Lag Technique	<i>41</i>
2 Achieving Fairness through Interactivity	<i>42</i>
3 Simulation Assessment	<i>44</i>
4 Results	<i>45</i>
4.1 <i>Interactivity and Fairness</i>	<i>47</i>
4.2 <i>About Scalability</i>	<i>49</i>
5 Wireless Home Scenario	51
1 Transmission Control Protocol	<i>52</i>
2 The IEEE 802.11	<i>52</i>
3 Queuing Delay	<i>53</i>
4 Proposed Solution	<i>54</i>
4.1 <i>IEEE 802.11 Parameters Setting</i>	<i>54</i>
4.2 <i>TCP Vegas</i>	<i>54</i>
4.3 <i>Limited Advertised Window</i>	<i>55</i>
5 Simulation Assessment	<i>58</i>
6 Experimental Results	<i>62</i>
6.1 <i>FTP Impact on Real-Time Entertainment Applications</i>	<i>62</i>
6.2 <i>Shadowing and Distance Impact on TCP Throughput</i>	<i>62</i>
6.3 <i>Appropriately Setting MAC Layer Parameters</i>	<i>65</i>
6.4 <i>Utilizing TCP Vegas in Place of TCP New Reno</i>	<i>68</i>
6.5 <i>Limiting TCP's Advertised Window</i>	<i>72</i>

6.6	<i>Summarizing Results</i>	75
6	Conclusion	77
1	Future Work	77
4.1	<i>Utilizing IEEE 802.11e</i>	78
4.2	<i>MMOG in Vehicular Networks</i>	78
	Acknowledgment	81
	References	82

List of Figures

Fig. 1	The exponential increase of MMOG subscriptions (1997-2005) [25]	9
Fig. 2	MMOG revenue by Region (2003) [25]	9
Fig. 3	Half-Life message format [171]	11
Fig. 4	Online game architectures	24
Fig. 5-a	Examples of obsolescence and correlation: event e_2 makes obsolete e_1	27
Fig. 5-b	Examples of obsolescence and correlation: event e_c is correlated to e_1 and rescinds the obsolescence of e_1	27
Fig. 6	Discarding probability functions for ILA-RIO	30
Fig. 7	ILA-RIO algorithm	30
Fig. 8	The adopted configuration	32
Fig. 9	Percentage of events with GTD over GIT; AIDT = 30ms	35
Fig. 10	Cumulative function of the GTDs in a scenario with 7 GSSs; AIDT = 30ms	35
Fig. 11	Percentage of discarded events; AIDT = 30ms	36
Fig. 12	# of activations of phase ON and phase 2 for ON-OFF and ILA respectively; AIDT = 30ms	36
Fig. 13	Average size of non-interactive bursts; AIDT = 30ms	37
Fig. 14	Total Number of bursts of non-interactive events; AIDT = 30ms	37
Fig. 15	Percentage of discarded events; AIDT = 45ms	37
Fig. 16	Percentage of discarded events; AIDT = 60ms	37
Fig. 17	Percentage of events with GTD over GIT; AIDT = 45ms	39
Fig. 18	Percentage of events with GTD over GIT; AIDT = 60ms	39
Fig. 19-a	Probability of obsolescence = 50%: event percentage having GTD > GIT	40
Fig. 19-b	Probability of obsolescence = 50%: percentage of discarded events	40
Fig. 20-a	Probability of obsolescence = 90%: event percentage having GTD > GIT	40
Fig. 20-b	Probability of obsolescence = 90%: percentage of discarded events	40
Fig. 21	Delay definitions	42
Fig. 22	Game servers deployment	44
Fig. 23	Interactivity and fairness improvement (left) and dropped events (right) with GIT=150ms and AIDT=30ms	45
Fig. 24	Interactivity and fairness improvement (left) and dropped events (right) with GIT=200ms and AIDT=30ms	46
Fig. 25	Interactivity and fairness improvement (left) and dropped events (right) with GIT=250ms and AIDT=30ms	46
Fig. 26	Interactivity and fairness improvement (left) and dropped events (right) with GIT=300ms and AIDT=30ms	46
Fig. 27	Interactivity and fairness improvement (left) and dropped events (right) with GIT=150ms and AIDT=20ms	47
Fig. 28	Interactivity and fairness improvement (left) and dropped events (right) with GIT=200ms and AIDT=20ms	47
Fig. 29	Interactivity and fairness improvement (left) and dropped events (right) with GIT=250ms and AIDT=20ms	48
Fig. 30	Interactivity and fairness improvement (left) and dropped events (right) with GIT=300ms and AIDT=20ms	48
Fig. 31	Interactivity and fairness improvement (left) and dropped events (right) with GIT=150ms and AIDT=10ms	49
Fig. 32	Interactivity and fairness improvement (left) and dropped events (right) with GIT=200ms and AIDT=10ms	49
Fig. 33	Interactivity and fairness improvement (left) and dropped events (right) with GIT=250ms and AIDT=10ms	50
Fig. 34	Interactivity and fairness improvement (left) and dropped events (right) with GIT=300ms and AIDT=10ms	50
Fig. 35	Pseudocode of TCP Vegas congestion control	55
Fig. 36	Comparison between regular and limited sending windows	56
Fig. 37	Simulated topology	57

Fig. 38	In-home wireless scenario.....	58
Fig. 39	Example of online gaming interarrival delays	60
Fig. 40	Example of online gaming jitter.....	60
Fig. 41	FTP total throughput with different user-AP distances; shadowing deviation = 9, MAC queue size = 50pkts	61
Fig. 42	FTP total throughput with different shadowing deviation values; user-AP distance = 10m, MAC queue size = 50pkts	61
Fig. 43	Example of TCP congestion window; max MAC retransmissions = 2.....	62
Fig. 44	Example of TCP congestion window; max MAC retransmissions = 3.....	63
Fig. 45	FTP total throughput with different MAC queue sizes; user-AP distance = 10m, shadowing deviation = 9	63
Fig. 46	FTP total throughput with different MAC queue sizes; user-AP distance = 5m, shadowing deviation = 9	64
Fig. 47	TCP congestion window; MAC max retransmissions = 4, buffer size = 50 packets	66
Fig. 48	Online game interarrival time; MAC max retransmissions = 4, buffer size = 50 packets.....	67
Fig. 49	Online game jitter; MAC max retransmissions = 4, buffer size = 50 packets.....	67
Fig. 50	TCP Vegas congestion window; $\alpha = 1, \beta = 3, \gamma = 2$	68
Fig. 51	Online game interarrival time with concurrent TCP Vegas; $\alpha = 1, \beta = 3, \gamma = 2$	68
Fig. 52	Online game jitter with concurrent TCP Vegas; $\alpha = 1, \beta = 3, \gamma = 2$	69
Fig. 53	Online game interarrival time with concurrent TCP Vegas: $\alpha = 5, \beta = 10, \gamma = 8$	69
Fig. 54	Online game jitter with concurrent TCP Vegas; $\alpha = 5, \beta = 10, \gamma = 8$	70
Fig. 55	TCP Vegas congestion window; $\alpha = 1, \beta = 3, \gamma = 2$	70
Fig. 56	Online game interarrival time with concurrent TCP Vegas: $\alpha = 3, \beta = 7, \gamma = 5$	71
Fig. 57	Online game jitter with concurrent TCP Vegas; $\alpha = 3, \beta = 7, \gamma = 5$	72
Fig. 58	Statistical values when employing limited advertised window.....	73
Fig. 59	TCP behavior with limited advertised window and $C = 18\text{Mbps}$	73
Fig. 60	Online game interarrival time with concurrent TCP, limited advertised window, $C = 18\text{Mbps}$	74
Fig. 61	Online game jitter with concurrent TCP, limited advertised window, $C = 18\text{Mbps}$	74
Fig. 62	Statistical values for the various schemes	75
Fig. 63	Hybrid architecture for distributed game entertainment in heterogeneous scenarios including car-networking.....	79

List of Tables

Table I	Sending GSSs involved in the simulations	33
Table II	Maximum, minimum, average and standard deviation of the GTDs (ms); AIDT = 30ms	36
Table III	Maximum, minimum, average and standard deviation of the GTDs (ms); AIDT = 45ms	38
Table IV	Maximum, minimum, average and standard deviation of the GTDs (ms); AIDT = 60ms	39
Table V	Percentage of obsolete and valid discarded events in ILA-RIO	40
Table VI	Simulation configuration of the wired links	57
Table VII	Simulated application flows	58
Table VIII	Changing parameters in the simulated configurations.....	59
Table IX	Gaming flow jitter statistics; max MAC retransmissions = 4, shadowing deviation = 9... ..	64
Table X	Gaming flow jitter statistics; max MAC retransmissions = 3, shadowing deviation = 9... ..	65

Chapter 1

Introduction

The current status of the Internet as a widely utilized tool and the overwhelming development of wireless access technology lead us to a future in which the synergy between wireless and the Internet will be integral part of our everyday life. Virtual libraries, remote-working, video-telephony and voice over IP, traffic control, remote-medicine, video and music on demand, on-line games, location based resource discovery, navigation support, are only a few of the innumerable services that will be ubiquitously available [5, 6, 9, 10, 12, 37, 40, 42, 45, 46].

People will be continuously connected during the whole day, regardless of their location and utilizing a plethora of traditional or new devices. Furthermore, connectivity will follow the *always best connected* paradigm so that the connection will seamlessly switch from one access technology to another guaranteeing always the best conditions for customers (i.e., cost, bandwidth, coverage, personal preferences) [13].

We are crossing a technology threshold that will revolutionize every area of our lives. It will affect all of our everyday habits and businesses in ways far more pervasive than people may imagine. Devices that we use today for a limited range of special purposes will become multiple application platforms. Even common objects such as wristwatches, cars, PDAs are evolving and their enhancement toward multipurpose tools will accelerate as we move forward.

Wristwatch capabilities will be augmented making it able to communicate, download/play music, keep personal/medical information, identify us to our car/home/devices, etc. Cars will be elevated from a simple transportation vehicle to an office on the move, as well as an information provider and entertainment center. Passengers will be allowed to access the Internet, engage in teleconferencing, play distributed videogames, learn location based information as low traffic paths to the destination or special offers for hotel reservations [114, 115, 116], participate in ad hoc or mesh networks [92], be part of an urban grid [93], etc.

Paper money and coins will be completely substituted by electronic transactions. Credit cards are already more and more frequently used by customers for their real or on-line purchases. The availability of always connected devices will further push commerce in this direction. Pocket sized PDAs, in fact, could be enhanced to be an easy-to-use way of payment. Customers will buy their objects and services connecting to a web page or simply making their PDAs interact with the cashier. Even tickets for events or travel will disappear in the paper format: PDAs or equivalent multipurpose devices will be able to provide the information required to authenticate the completed transaction.

As a proof for this forthcoming scenario, hot spots are rapidly increasing in number providing people with wireless connectivity in almost all the buildings they enter (e.g. home, work, cafeteria, etc.) and all the streets they walk and drive in. A large amount of static and nomadic users relies today on a wireless access to the Internet to run their favourite applications: email, web surfing, P2P file sharing, chatting, and video/audio downloading/streaming.

Mobile and highly mobile users will soon be using both traditional and innovative applications. Novel services will be provided whose utility sensibly rises for mobile users: location-dependent information, information exchange with people around, mobile market, safe driving alert system, urban grid for traffic control, Video/Voice over IP, Text/Voice-Chat on the move, and many others that we are still not able to imagine [16, 22, 23, 47, 82, 113].

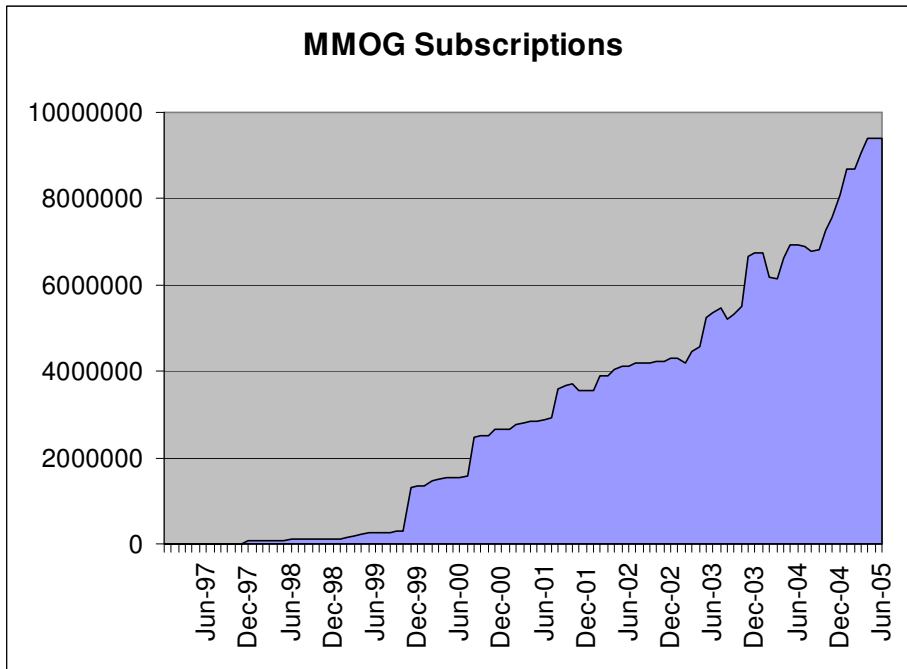


Figure 1. The exponential increase of MMOG subscriptions (1997-2005) [25].

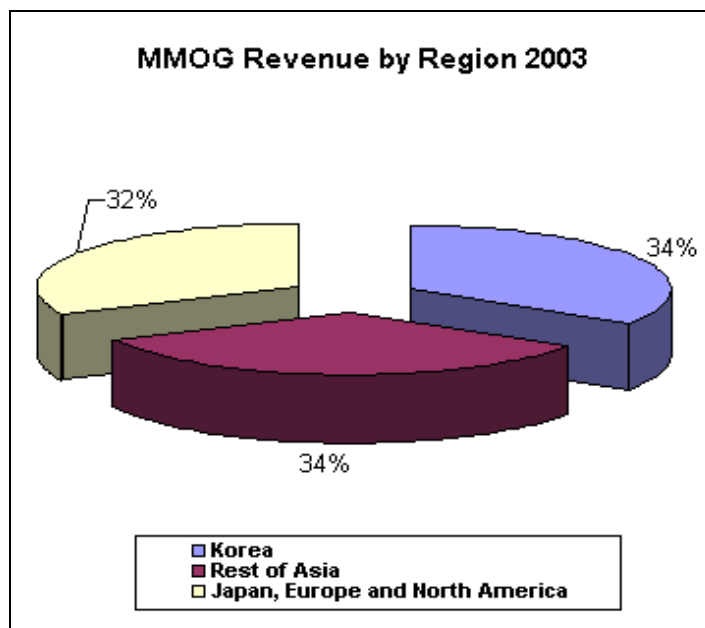


Figure 2. MMOG revenue by Region (2003) [25].

Houses will not be spared by this *wireless revolution*. Nowadays, home networking is still limited to few PCs and electronic equipments located in different rooms of the house; yet the vision for the near future includes many devices networked within a single household and connected to the Internet [117]. The *domotic philosophy* is steadily increasing the number of habitations which combine technology and services to improve living in the areas of safety, comfort and technical management. *Smart houses* will

be endowed with devices, appliances and sensors, communicating through the means of wireless technologies and embodying active partners in managing our everyday life.

The trend of the market strongly suggests that entertainment applications are going to play a major role in this scenario and, among them, video games are gaining more and more attention. In the last decade, in fact, thanks to their impressive progression in plunging players into terrifically realistic and capturing virtual worlds, videogames have expanded their market with a persistent and accelerating growth (see Fig. 1) [24, 25].

Furthermore, this market is still far from being mature and presents large growth margins. In fact, as can be seen in Fig. 2, a very large portion of MMOG revenues comes exclusively from South Korea (~50 millions of habitants). By projecting this value over the whole world and pretending that all the other developed country will eventually reach the same percentage of online players on their whole population, we can have a clue of the astonishing potentiality which hides behind this market.

Nowadays, two main reasons above the others attract an increasing number of researchers and practitioners toward the MMOG field. The first one is the explosive growth of the computer games market, with an increasing trend whose end is still not in sight [26]. The second reason is represented by the correlation between problems that emerge in developing innovative game experiences and those typical of various other conventional research fields in Computer Science and Engineering. Indeed, creating enjoyable online games requires the convergence of solutions belonging to extremely diverse technical areas. Examples are represented by networking, computer graphics, animation, music and sound, multimedia design, Artificial Intelligence (AI), human-computer interaction, software engineering, virtual environments and distributed simulation [27, 28, 29, 30, 31, 32].

Massively Multiplayer Online Games (MMOGs) are further extending the boundaries of what has been defined “the tenth art” with the possibility of contemporary engaging, in the same virtual scenario, millions of players located all over the world. Indeed, one of the main elements that determine the success of a game has always been represented by the possibility to engage in multiplayer sessions. Humans are social beings and generally consider challenging other humans as funnier than playing alone against an artificial intelligence. Nowadays, the boom in the Internet usage has brought the logistic advantage of an always available virtual arena where millions of players can contemporary participate in electronic multiuser amusements.

Furthermore, the exploding market of connectable handheld devices, always looking for new killer applications, pushes the game industry to propose effective distributed game platforms proficient at engaging an unlimited number of contemporary users [33, 34]. This large and emerging market is driving researchers and practitioners to develop novel distributed solutions able to efficiently sustain interactive multiplayer networked game sessions over best effort networks such as the Internet [7, 35, 36, 38].

1 Problem Statement and Thesis Contribution

In this Thesis we focus on problems that arise in heterogeneous networks that involve also wireless users when trying to deploy a highly interactive entertainment system; in this context, MMOGs represents an emblematic and challenging example.

We define the *game state* in the system as the set of information that univocally describes the current configuration of the game. Players perform moves based on their perception of the game state. Specifically, their terminals periodically render a projection of the game state on screens. Players can hence be aware of the surrounding virtual world and influence it by generating *game events* which includes, but are not limited to, the movement of an avatar, hitting/missing a target, changing difficulty level.

In its simplest version, with no particular techniques employed, Game events are propagated from clients to some *decision point(s)* which utilize(s) them to determine the new game state that will be broadcasted. Decision point(s) could be represented by a centralized server, or a constellation of mirrored servers, or just other players, depending on the underlying architecture (see Section 2.5). Indeed, it is of particular interest to discuss potentially efficient game server architectures by analyzing network and computational issues related to the maintenance of a consistent game state in the whole game platform.

```

typedef struct usercmd_s
{
    // Interpolation time on client
    short lerp_msec;
    // Duration in ms of command
    byte msec;
    // Command view angles.
    vec3_t viewangles;
    // intended velocities
    // Forward velocity.
    float forwardmove;
    // Sideways velocity.
    float sidemove;
    // Upward velocity.
    float upmove;
    // Attack buttons
    unsigned short buttons;
    //
    // Additional fields omitted...
    //
} usercmd_t;

```

Figure 3. Half-Life message format [171].

Residential broadband connectivity is currently becoming more and more common. Moreover, MMOGs generally utilizes very small packets thus easing the bandwidth requirement for this kind of application. As a confirmation, Fig. 3 shows the format of packets sent from client to server in Half-Life [74, 171]. Fields represent basic information that could be contained within few tens of bytes [146].

In this context, we intend to provide cutting-edge MMOG experience both to wired and wireless users. To this aim, we propose a scalable system able to support fairness and interactivity even when players are connected through the wireless medium. More in detail, in order to ensure optimal performance to players we have to split the problem into two sub-parts which requires specific solutions applied by different subjects.

The first sub-part regards the communications and synchronization among game servers and represents the portion of the total connectivity that can be handled by the MMOG service provider (either directly, or through ISP-domain managers). It is up to the MMOG service provider to deploy the ultimate MMOG technology in this portion of the game platform.

The second sub-part, instead, is concerned with the links between game servers and their engaged players, thus including also the wireless hops. This portion of the connectivity is out of the control of the MMOG service provider: customers decide on their own whether they are going to subscribe cable/DSL connectivity, set up a wireless last hop, and/or install a particular Media Center. Yet, it is a critical part, as it generally includes the bottleneck of the connection and may include large queuing delays.

Therefore, to obtain an optimal solution we propose to proceed through successive steps and address the two sub-problems one independently from the other. Our mechanisms complement each other since their scopes are detached (even if connected) and, although they generate the best performance result when combined, they produce benefits also when singularly applied.

In particular, for the first part we exploit a hybrid architecture combining both the advantages of client-server and peer-to-peer paradigms. Our solution deploys over the network a constellation of communicating replicated Game State Servers (GSSs), each of which locally maintains a redundant version of the game state. Each GSS manages and updates its copy of the game state as follows: i) it collects game events coming both from its engaged players and from other GSS peers; ii) it forwards to all other GSSs the events generated by its connected players; iii) it updates the game state considering the set of received game events; iv) it finally delivers the newly updated game state to its connected players.

It goes without saying that, within this scenario, an efficient event synchronization scheme among GSSs needs to be employed to guarantee a consistent and responsive evolution of the game state. Indeed, one of the key factors in determining the success of an online game is represented by the ability to rapidly deliver events among the various GSSs. While simpler turn-based games do not have to face this problem, as only one player is allowed to perform an action at any given time, the task of providing players with responsiveness and real-time interactions is probably the most stringent requirement for MMOGs (especially for fast-paced MMOGs). In fact, in case of intense traffic in the network or when excessive computational loads are slowing down some GSSs, the game delivery activity turns out to be quite complex. As a consequence, the responsiveness of the distributed game system may be jeopardized.

With this in mind, we propose a scheme that takes inspiration from Active Queue Management techniques to maintain the game event delivery delays under a human perceptivity threshold and uplift the playability degree of MMOGs. Our scheme follows a holistic approach aiming at preserving also consistency and fairness. At the basis of our scheme lies the idea of exploiting the semantics of the game and, in particular, the notion of obsolescence. Simply put, obsolescence entails that during a game session some events can lose their significance as time passes, i.e., new actions may make the previous ones irrelevant. For example, where there is a rapid succession of movements performed by a single agent in a virtual world, the event representing the last destination supersedes the previous ones (obsolete events).

Obsolescence allows the system to drop those game events that lose their importance during the game evolution. Discarding superseded events for processing fresher ones may be of great help for delay-affected GSSs. This means that, during the game event exchange activity, while responsive GSSs may deliver all game events to provide their connected clients with a fluent game state evolution, those GSSs that are experiencing loss of interactivity may skip the execution of obsolete events in order to speed up the event processing activity, thus gaining interactivity.

Upholding interactivity may be useful also to the aim of ensuring fairness. In fact, we demonstrate how our system is able to take advantage of the reduced transmission time to magnify the efficiency of a local lag-type algorithm in ensuring fairness without compromising interactivity. This represents a very important result since it contradicts the general belief that interactivity and fairness embodies antithetic objectives.

However, even if our scheme is proficient in maintaining a high degree of responsiveness among game servers, still problems may arise at the edges of the considered topology. This represents the aforementioned second part of our problem.

Concurrent traffic may generate queues that build up at the last (or first) link of the connection, thus delaying the game event delivery. This problem is worsened in case of players relying on wireless connectivity. The wireless medium, in fact, is naturally prone to be easily shared by several contemporary users who may interfere with each other. The applications run in this context may vary and we demonstrate as some of these may be particularly harmful toward real time traffic (online games but also video-streaming, video-chats, etc.). In particular, we show how the very popular TCP-based FTP application for downloading files increases queuing delays to such an extent that responsiveness may be completely jeopardized.

To this aim, we propose and compare some possible solutions aimed at maintaining low transmission delays even on the client-server link. In particular, we evaluate different approaches which involve diverse networking layers: i) appropriately setting IEEE 802.11b/g MAC layer parameters, ii) employing

TCP Vegas in place of the traditional TCP New Reno and, iii) limiting the advertised window for TCP flows. Extensive simulation results for realistic wireless scenarios are provided.

The contribution of this Thesis is hence that of providing a scalable MMOG architecture endowed with a set of solutions able to ensure interactive, fair, and consistent game event delivery even for frenetic action games. Summarizing it with a motto: *fast and fair event delivery for fast and furious online games*.

2 Thesis Outline

The remaining of this Thesis is organized as illustrated below.

In Chapter 2 we introduce the reader in the MMOG field by analyzing the most important issues in this research area. In particular, we overview key requirements and fundamental problems shared with traditional field in computer science, we summarize previous work that constitutes helpful background, we propose a model for MMOG delays, and we analyze possible system architectures.

Chapter 3 depicts the general framework that we use to guarantee fast synchronization among servers. It introduces the concepts of obsolescence and correlation and explains how to use them in combination with queuing management techniques in order to uplift the interactivity level of the system. Extensive results are provided to show the benefits attainable through our scheme.

Our proposed synchronization scheme is further refined in Chapter 4 where we show how increasing interactivity is not incompatible with aiming at fairness among players. On the contrary, the latter could be more easily achieved through the means of the former. We present experiments and outcomes that prove our assertion.

In Chapter 5 we describe the wireless home as an important scenario where MMOGs are going to play a major role. In particular, we demonstrate how MMOG performance could be affected by concurrent traffic of diverse nature. We then propose and evaluate different possible techniques to find an efficient and effectively deployable solution.

Finally, Chapter 6 concludes this Thesis and proposes a couple of very promising directions for future research work.

Chapter 2

Background

We present a general introduction to the MMOG research field. In particular, we first discuss the major MMOG issues and wrong assumptions generally made by online game developers. We then explain fundamental problems that MMOGs share with other traditional field in computer science, thus elevating themselves to the rank of scientific topic. After this, we propose a model for delays in MMOGs. Finally, we analyze possible architecture for online games and survey related work in this area.

1 MMOG Issues

Several elements contribute to the success of a game: capturing graphics and visual effects, original plot, fair difficulty, multiplayer sessions, high interactivity level, etc. MMOG is, by definition, a class of games devised to engage a multitude of users, contemporary present in the same virtual arena even if physically located far away from each other. This scenario could be realized through the Internet; however, the best effort nature of the Net poses several challenges before being able to proficiently deploy really scalable and interactive game sessions.

In this sense, a MMOG is subject to similar major issues and fundamental principles which emerge when developing a general distributed application [75]. These are well summarized by Peter Deutsch's "*Eight Fallacies of Distributed Computing*" [130].

Essentially, developers building a general distributed application (or a MMOG) may fall in one or more of the following wrong assumptions generating a system that is intrinsically unable to guarantee high performance:

1. "*The network is reliable*". Errors, losses, and wrong order delivery are common issues when utilizing a network to transmit messages. The larger is the scale of our MMOG, the higher is the probability that one of these events happens contradicting this assumption. Therefore, the application should be able to face these issues through mechanisms able to provide reliability without affecting the interactivity of the game system.
2. "*Latency is zero*". Depending on the topology of the game network and on its conditions, delays may conspicuously increase. Objective delays as perceived by players are impacted both by network delays and by processing workload delays. At the same time, since MMOG requirements for interactivity and fairness, latency is the main issue to address in this kind of application.
3. "*Bandwidth is infinite*". In a LAN scale network game this assumption may still hold since the available bandwidth is generally much larger than the required one to run a MMOG. On the other hand, when moving to wide area networks, congestion due to a very large and unpredictable traffic may generate bottlenecks that could invalidate this assumption.
4. "*The network is secure*". Security in a large scale MMOG includes various issues that can be shared with traditional distributed systems or even be completely new. Authentication, subscription transactions, and cheating represent the most relevant among them.

5. “*Topology doesn’t change*”. In a LAN, failures can suddenly lead to having isolated PCs that cannot anymore communicate among themselves nor, obviously, be engaged in a MMOG. Considering the larger scale case of the Internet, failures are addressed by the presence of redundant alternatives. Yet, the consequent topology change and congestion increase may noticeably impact on the system performance.
6. “*There is one administrator*”. MMOGs can be deployed across large geographic areas requiring a decentralized control of networking resources. Therefore, multiple domains and administrators may be involved in the resource allocation.
7. “*Transport cost is zero*”. The exchange of messages among nodes is a characteristic that belongs to the nature of a MMOG. These messages have to be forwarded through some infrastructure that needs to be provided and therefore bought/rented. Moreover, as the involved network spreads across different domains, resources among shared infrastructures have to be allocated to support the MMOG. Transport costs are hence inflated by providers’ rates.
8. “*The network is homogeneous*”. The Internet is heterogeneous in nature and this characteristic is going to become even more evident in future. As new wireless and mobile technologies are coming into the picture, different network accesses, as well as different devices, are offered to players. Moreover, the core of the Internet evolves gradually and slowly since the cost and the large number of domains involved. Evolution at its edge, instead, proceeds by leaps and bounces, thus generating an interoperating mix of components that increases the heterogeneity of the infrastructure.

As a demonstration of the wide diffusion of these wrong assumptions, many games have been initially developed to be singularly played on PCs. Since more and more houses have become endowed with high speed connectivity such as cable and DSL, game vendors started to endow their products with the possibility of playing online with other users. However, this feature has been added as an extension of single use designed games, thus clearly inheriting the above wrong assumptions. Instead, MMOGs aimed at providing an amusing experience to their players have to be designed since the beginning, as intended to be played over the Internet and taking into account all the involved implications.

2 Fundamental Problems

Delving into the Internet, the first video game appears to be a simple *Tennis for Two* created by Higinbotham in 1958 to entertain visitors of the Brookhaven National Laboratory, a US nuclear research lab in Upton, New York. Since then, video games have evolved from a simple pastime into a real business and an important research field.

Indeed, certain games simulate, relatively cheaply and safely, situations which could be much more expensive or dangerous in the real world. Technical solutions introduced by game developers are now employed in medical surgery, military simulations, distributed simulations, virtual reality, interactive collaboration, distance learning, e-commerce and manufacturing systems. In this sense, several works have been presented that describe the convergence between game technology and non-game applications.

For example, interactive storytelling may gain benefits by employing synchronization schemes developed for online games [64, 125]. Collaborative applications require augmented reality interfaces to support group interaction [124]. Combat video games are utilized to enhance strategic, combat, and decision-making skills of military commanders [126, 129, 178]. Psychology is requesting for immersive collaborative virtual environments to investigate human behaviors and interaction [122, 123], as well as to support traumatized patients [121]. Real time video-streaming, voice over IP, and online games are all extremely delay sensitive applications that need fast delivery support [40, 99]. Finally, surgery simulations and video games share a quest for realistic object behavior, immediate response to given commands, and high-quality images [128].

These represent only a few of the innumerable intersections between video games and non-game applications. It is hence easy to see how issues emerging with online games represent fundamental

problems also shared by other traditional fields in computer science, and how adapting elements of computer games may enable the creation of compelling user experiences in several domains [120].

Under a networking point of view, distributed multiplayer games are characterized by four main requirements which are intrinsically correlated and correspond to major research challenges involved with MMOG, namely: interactivity, consistency, fairness, and scalability.

Interactivity (or *responsiveness*) refers to the delay between the generation of a game event in a node and the time at which other nodes become aware of that event. Therefore, it includes both the network latency and the processing time. Having a high level of interactivity represents a fundamental quality for a MMOG. In order to assure an enjoyable playability to the final user, external stimuli generated by players need to be processed under a human-perceptivity threshold. This means that the time elapsed from the game event generation at a certain node and its processing time at every other node participating in the same game session must maintain a low average value. Unfortunately, not only could this be very hard to be accomplished in a best-effort network, but we will probably face also a high variance in the delivery time of game packets. Variable congestion conditions in Internet could, in fact, result in sudden slow down of the perceived game fluency on screen. Moreover, players in the same virtual arena can increase in number, even sensibly, with almost no predictability. Some game server may thus experience impulsive computational load and loose interactivity. These problems are obviously amplified when plunged into a wireless scenario.

Consistency, on the other hand, regards the contemporary uniformity of the game state view in all the nodes belonging to the system. Depending on the features of the game, consistency requirements may be absolute or partial. The easiest way to guarantee absolute consistency would be that of making the game proceed through discrete locksteps [98]. At each step, the system waits until having received all the actions generated by the final users; only at this moment a new instance of the game is produced and propagated to all the nodes. Having a single move allowed for each player and synchronizing all the agents before moving toward the next round, for sure grants absolute consistency but, on the other hand, impairs the responsiveness of the system. Obtaining both absolute consistency and high interactivity would require the employment of almost unlimited network and computation resources (very high bandwidth, very low latencies, very high speed at server to process events). A trade-off between the two attributes needs to be found in order to develop a proficient game platform.

Fairness among users is another major issue in MMOGs. In fact, to ensure a rewarding game experience to users, every player needs to possess the same chances of winning than any other, regardless of the different network connections. Indeed, we are here interested in networking fairness for MMOGs as it represents one of the major issues that need to be addressed when developing new online games. In this context, relative delays have to be considered as important as absolute ones. Simultaneous game evolution with identical speed should be guaranteed as much as possible to all the participants [78]. To this aim, introducing appropriate artificial delays before displaying both generated and received game events may represent a feasible solution. However, similar to consistency, aiming at a full fairness may result in excessively increasing game delays, thus jeopardizing interactivity. Indeed, it is generally believed that interactivity and fairness/consistency embodies antithetic requirements [104, 105, 127].

Scalability regards the capability of the system to provide an efficient support to a large community of players. Regarding this point, it should be noticed that the interest of companies in online gaming emerges from the huge revenues that may be generated by a very elevated number of customers. Besides, humans are social beings which enjoy the presence of others in most of their amusement activities (i.e. team sports, movies in theatres) and the competition in challenging their skills against real adversaries. However, especially in the case of fast-paced MMOG, scalability is sometimes sacrificed to maintain a high degree of interactivity. In some cases, in fact, the system could deny the MMOG access to some users depending on their experienced delays [179]. These delays could have been generated by several factors such as the location of these users with respect to the server, the network conditions, and the current computational load of the system. Limiting the access to some customers obviously eases the achievement of interactivity, consistency, and fairness, but at the cost of a reduced scalability.

A well designed game architecture could help in devising a MMOG which possesses all the required qualities. The positioning of servers affects network latencies; thereby, they should be optimally located to efficiently serve their customers [76]. Moreover, every time a new scheme is proposed as a solution for MMOGs, all the four aforementioned key factors should be ensured and verified. Generalizing this concept, developers should follow a holistic approach when designing a new MMOG, considering the

whole set of requirements and aiming at the intersection of their solutions. Addressing only one requirement, in fact, could produce the unexpected and undesired result of jeopardizing the others.

3 Related Work

Here, we present a commented summary of research contributions that have to be considered in our scenario. For each of them, we introduce the proposed technique and then highlight both the advantages and disadvantages of using it. In particular, we first evaluate existing techniques aimed at increasing the responsiveness degree of online games. We then analyze solutions proposed to ensure fairness and consistency (solutions developed to satisfy the fairness can generally be extended to satisfy consistency too, and viceversa). For the sake of completeness, we also provide a panoramic of anti-cheating mechanisms analyzing their impact on the other requirements. Finally, we are interested in efficiently extending the boundaries of MMOGs also to the wireless domain. Therefore, we discuss research about the impact of a wireless link on measured performance.

3.1 Responsiveness

Trying to improve the responsiveness of a distributed game architecture, two main causes for delays have to be analyzed: network latencies and computational costs. Several research works have already brought contributions to the factual developing of efficient synchronization schemes. *Compression* and *aggregation* consider networking having a dominant position when dealing with the delays and thus with the playability of a MMOG [39].

More in detail, packet compression tries to speed up transmissions by reducing the amount of payload bytes to be transmitted. However, the provided benefits are very limited as the size of online game packets is already very small (see Fig. 3 and [146]). Aggregation, instead, merges packets together in the attempt of reducing the transmission overhead. Both compression and aggregation pay the achieved latency benefits with an increment in computational costs. Moreover, aggregation can generate further waste of time if a transmission is delayed while waiting for having available other events that could be aggregated.

In the attempt of reducing both the traffic load in the network and the computational cost to process each game event, *interest management* techniques rely on the *area-of-interest* concept to filter events. Players, in fact, are not impacted by game events happening far away from their current virtual position. Therefore, sending all the game events to all the players represents a waste of network and computational resources with no utility counterpart. Instead, an appropriate multicast based scheme could match every packet with the nodes that really need to receive it in order to reduce both the traffic on the channel and the processing burden at each node.

On the other hand, a tradeoff exists between the computation spared at the destination by receiving only a limited number of packets and the required one for implementing the filtering scheme at the sending GSS. Moreover, interest management techniques could further delay packet delivery when applied to games having almost all the game actions to be forwarded to the majority of the participants. Finally, area-of-interest techniques do not help in reducing latency as there is no correspondence between players' position in the virtual world and their actual location in the real one. In simple words, players positioned in the same virtual area-of-interest might be physically located very far from each other, thus not gaining any benefits in terms of propagation and queuing delay.

Slightly detaching playability from the real responsiveness of the network, *optimistic* algorithms for synchronizing game state at servers can be utilized in order to avoid delay perception at destination. In case of high delays in events forwarding between GSSs, in fact, an optimistic approach may execute game actions on clients before really knowing if ordering would require processing other on-the-way ones first. Game instances are thus processed without wasting any time in waiting for other packets that might arrive later. On the other hand, this performance gain is paid with some occurrence of temporary consistency loss. Standard *Time Warp* and *Breathing Time Warp* represent typical exemplars of this family of algorithms [52, 53, 105]. *Rollback* based techniques are exploited to reestablish the consistency of the game state but may further impact on the responsiveness of the system.

Based on this concept, a new synchronization mechanism for online games called *trailing state synchronization* was presented by Cronin *et al.* [59]. According to this approach, every GSS locally maintains a fixed number of copies of the game state, each of which is kept at a different simulation

time. In essence, each copy of the game state is associated with a particular execution, and each execution is delayed for a fixed time interval. If no inconsistencies are detected, the game proceeds in rendering the copy of the game state which is more forward in time; as a result, players perceive the game evolution as being very responsive.

With this trailing state synchronization, inconsistencies are identified by comparing the leading state (that optimistically processes game events without any additional delay) with the game states of the delayed executions that reorder and then process the received game events. If an inconsistency is detected, a rollback is performed by copying the (consistent) game state from the delayed execution to the leading execution and then the rolled back game events are re-processed in the proper order. It goes without saying that a tradeoff relationship exists between the number of game state copies maintained by each GSS and the number of game events that need to be re-processed in case of a rollback. Moreover, this scheme does not avoid short-term inconsistencies: it just provides a mechanism to detect and correct them after they happened.

Optimistic algorithms can be employed also at the boundaries of the game platform. Clients could, in fact, be equipped with the intelligence required to perform predictions of other players' movements. Obviously, unless utilizing a fully distributed architecture, servers have still to validate clients' predictions. If the prediction results wrong, the server(s) will eventually correct the mistake and restore a consistent game state. However, a full Round Trip Time (RTT) may elapse between the generation of an inconsistency due to wrong prediction and its correction. If this period of time surpasses the human perceptivity threshold, players could be annoyed by it.

In general, prediction techniques rest upon the assumption that a consistent game state exists at some point. From this starting point, new actions are predicted and executed on clients' screen. For instance, 50 frames per second represent a typical refreshing rate for an online game and correspond to one new game state sent from the server to the client every 20ms. Considering a connection of 100ms of RTT, this amounts to 5 game events that could be subject to rollback if the server does not confirm the correctness of the predictions [171].

A tradeoff between the speed of executing game actions and the frequency of short-term inconsistencies is hence present in this kind of solutions. Moreover, client side prediction techniques require that clients and servers share the same predictive code; this solution could hence be not applicable on clients equipped with low computational resources. Finally, having portions of code shared between all clients and servers gives an advantage to malicious players. Cheaters, in fact, could become aware of the next expected (and predicted) actions of regular players and modify their behaviors in order to obtain the maximum benefits out of it.

Extrapolation and *Interpolation* are particular instances of the optimistic approach. With extrapolation, the next action and position of a virtual object (even an avatar) are predicted by exploiting the current available information (i.e. current position, speed, direction). Actions can thus be rendered on the screen before really receiving the related game event.

In games where physical rules of the real world are respected, the task of predicting events forward in time is made easier. Unfortunately, most of the games available in commerce present very poor similarities with real physical laws or ballistic models. Indeed, virtual object in games are usually allowed to perform sudden actions featured with huge acceleration/deceleration and instantaneous change of direction. These unrealistic movements cannot be anticipated by extrapolation and result in inconsistencies that affect the player's perception of the game evolution. Depending on the extension of the extrapolated time, these errors can be more or less evident. Therefore a tradeoff exists between the impact of these inconsistencies and the augment of responsiveness that can be simulated by this scheme.

Interpolation is another technique that tries to guess the position/action of a virtual object at a certain moment without having precise information about it. In particular, interpolation uses exact information about two valid positions/actions of an object at two different moments to determine how to display the in-between movements of that object.

When interpolation is employed the constraint for reliability in game event delivery can be relaxed. For instance, if some of the movements of an avatar, leading from an initial to a final position, are lost, the system is still able to reconstruct the whole sequence. However, some critical movements cannot be interpolated without creating inconsistencies, thus requiring reliable transmission (see the notion of correlated events in Section 3.2).

3.2 Consistency and Fairness

The simplest way to enforce fairness and consistency in the game is that of proceeding through discrete *locksteps* [98]. Simply stated, the game evolves by marching in step and players have to wait their turn before making any action. Every game event is thereby reliably received by all the players before any new move could be instantiated. Consistency and fairness are hence ensured by the fact that before proceeding to any new step, all the nodes in the game platform will share the same identical game state view.

Unfortunately, this scheme is affected by some important drawbacks. The most relevant one is the fact that locksteps cannot be used for interactive games since updates are displayed on screens only after having received game events from all the players. This requires reliable transmission of game events and may sometimes require many seconds, thus making the game proceed very slowly. Moreover, if one node in the system fails, then all the others could wait indefinitely before receiving updates from that node and impeding any progress in the game evolution.

In [71, 105, 131], *Mauve et al.* presented an optimistic approach for the consistency control in networked multiplayer games. Their approach utilizes the local lag control mechanism combined with a modified Time Warp (executed only when necessary) in the attempt of solving the trade-off relationship existing between responsiveness and consistency.

In particular, *Mauve et al.* devised a synchronization approach based on the idea of intentionally decreasing the responsiveness of the application in order to eliminate short-term inconsistencies. Exploiting a local lag approach, game events are delayed for a certain amount of time before being executed by clients. This additional delay permits the reordering of the received events in order to minimize inconsistencies. However, the presence of this delay is not always sufficient to prevent short-term inconsistencies as game events might still arrive late due to jitter or packet loss in the network.

In [137], *Li et al.* presented a continuous consistency control mechanism for supporting networked multiplayer games. Similar to [105], this approach rests on the idea that game events should not be considered as discrete updates of the game state, but rather as continuous actions in the game world. In this context, the authors proposed a relaxed time-dependent consistency control scheme which gradually synchronizes the nodes in the system and ensures that the possible discrepancy among game states stored at each node never exceeds a predefined threshold.

As it is evident, this scheme does not avoid short-term inconsistency. Rather, it deals with them after a certain period of time to avoid that the divergence among the various game state views does not become excessive. However, players' gaming experience could still result affected.

Other approaches simply assume that delaying the game event processing activity for a predetermined amount of time may be sufficient to guarantee a uniform evolution of the game state at different nodes without any need to resort to rollback. Following this idea, *Diot and Gautier* depicted in [132] a synchronization mechanism which embodies an optimistic version of the well known conservative *bucket synchronization* algorithm.

Their approach assumes that there exists a processing deadline (which defines a time bucket) and that a correct evolution of the game requires that all game events are received prior to this deadline. In essence, the idea behind this scheme is that of ordering and processing game events at the end of the time bucket. If some game events are not received before the time bucket expiration, *dead reckoning* techniques are exploited to compensate event losses [134, 135]. A similar approach was presented also in [133]; in both cases, the main drawback is that dead reckoning does not ensure the full consistency of the distributed game state.

Short-term inconsistencies may still arise if the bucket size is set excessively small but, on the other hand, using large buckets may induce severe responsiveness degradation. Further, a complex problem is that of fitting the bucket size with the unstable condition of the Internet where large and variable jitter values may be experienced.

Fairness among users is a very important issue in online games. Every player has to possess the same chances of winning than any other. To this aim, it is fundamental that delays impact on all participants in equal measure. Fairness could be enforced by compensating latency differences among players with an appropriate queuing of game events at each client before presenting them onto the screen [54, 78, 94, 104, 105]. This artificial injection of an appropriate amount of delay is usually referred to as *Local Lag* technique.

In essence, this scheme normalizes the game state so as to have all the players virtually brought back at the time when a certain game event was generated. Unfortunately, pure Local Lag algorithm does not

solve the tradeoff between interactivity and fairness anticipated in Section 2.2.. Lags, in fact, can be inflated only up to the delay perceptivity threshold; otherwise all the players will be affected by excessive delays which jeopardize the interactivity degree of the whole game platform.

Finally, aiming at ensuring consistency, *Knutsson et al.* propose in [136] the use of coordinators to solve update conflicts that may occur in a peer-to-peer architecture. In particular, they split the game state management into different classes handled by different coordinators within a group of interested peers. Aiming at guaranteeing fault tolerance, they also propose the use of a primary-backup protocol to address possible fail-stop failures of coordinators. However, several issues arise in this context such as, coordinator election, fairness maintenance among different kind of nodes, authentication, and cheating avoidance from malicious intruders that become coordinators.

3.3 Cheating

A deep analysis on cheating and proposed solutions is out of the scope of this Thesis. Nonetheless, for the sake of completeness we provide a short panoramic of work on this topic. Moreover, research in this area generally proposes protocols which also include solutions aimed at sustaining other key requirements such as interactivity, consistency, and fairness. Conversely, none of the schemes reviewed till now proposes any cheat-proof mechanism.

A general taxonomy of cheating in online games can be found in [169]. A more detailed one, adopting a classification based on networking layers and then further categorizing the *protocol level cheats* class, is presented in [167]. In [167], the authors also present NEO, a low-latency event ordering protocol for a distributed architecture. Designed to improve the responsiveness of the game while preventing protocol level cheats, NEO takes inspiration from bucket synchronization and divides the time into equal intervals called *rounds*. Within these rounds, players send encrypted updates to each other and, in the following round, players send key to decrypt the updates sent. After this, a voting mechanism is employed to achieve *majority based* consistency and responsiveness.

Specifically, every new game state update is accepted by players only if a majority of them has received it on time. Consistency is hence maintained through a distributed voting system which collects positive votes from players who have received the game state update in time, and negative votes from delay affected players. In order to avoid interactivity loss, NEO does not wait for receiving all the votes. The relationship between game progression and communication reliability is weakened, any non-received votes at the end of the round are considered *abstentions*, and only the receiving of a majority of votes is required to consider the election as valid.

Therefore, NEO does not guarantee continuous global consistency or interactivity. Rather, its aim is that of ensuring these properties for a majority of players, and utilizing dead reckoning techniques in order to adjust the game state view for the rest of them (a minority). Obviously, depending on the definitions of majority and minority a different tradeoff can be found for interactivity and consistency. However, the length of a round is limited by the maximum latency acceptable to still consider the system as responsive, and no effort is devoted to facilitate a more efficient solution for this tradeoff by effectively decreasing the delays caused by traffic conditions. On the contrary, a bucket synchronization mechanism which also requires multiple rounds in order to send updates, encrypt them, and voting for determining their consistency, definitely affects the responsiveness of the system, and is therefore not well suited for highly interactive games.

An extreme solution for the protocol level cheats makes use of the aforementioned concept of locksteps. In particular, *Baughman and Levine* design in [168] a protocol that forces players to reveal their moves before actually performing them and, most important, before receiving any plain-text move from any other player. Thereby, this scheme prevents anyone from performing late changes of their moves based on the knowledge about other players' actions.

Two steps characterize this protocol. First, each player has to communicate her/his next move utilizing an encrypted message; the move is hence *committed* even if not revealed ahead of time. Finally, each player sends her/his move in clear. The authors also propose the utilization of area-of-interest techniques to reduce the overhead in the system caused by the reliable transport of commitment messages. Nonetheless, any lockstep based mechanism leaves unresolved several problems and, among the others the most critical one is the lack of responsiveness (as already seen in Section 2.3.1).

Another method based on locksteps is presented in [170]. In this work, *Chen and Maheswaran* developed a mechanism for P2P architecture composed by two protocols: the first one aimed at ensuring

fairness among all players regardless of their latencies, and the second one able to prevent certain types of time-cheats.

The fairness ensuring protocol makes use of specific nodes called *pulser*, and of network *sensors*. The former are elected nodes which periodically broadcast game state updates (*pulses*), while the latter detect the current status of the network and adapt the pulses sending rate in order to meet the fairness requirement.

Pulsers are also used to implement encryption on game messages in order to avoid time-cheats. In particular, they forward encrypted messages at a pace which corresponds to that one achieved by the slowest player. The authors demonstrate how this is sufficient to prevent faster players from maliciously peek into the future with respect to slower players.

Unfortunately this scheme presents several lacks. First, there is no real guarantee about the fine precision of the estimated network conditions; this is particularly true in the case where nodes are dispersed and sensors are not in close proximity of each of them. Second, this scheme works only if the whole structure (i.e., sensors and pulsars) can be adequately distributed over the Internet. Indeed, providing efficient clusterization of the nodes with respect to the sensors and the pulsers remains an open problem for this scheme. Finally, the anti-cheating mechanism only addresses time-cheats, while many other possible ways to gain unfair advantages have not been treated.

3.4 Performance Evaluation over Wireless Links

In recent years, many researchers have focused their studies on the problems encountered in a wireless environment [15, 56]. We try to limit the scope of our survey to those works that are related to problems we are trying to address in this Thesis and to the solutions we propose and evaluate in Chapter 5.

Focusing on MAC layer retransmissions, TCP and UDP flows have been tested by *Nam et al* over a IEEE 802.11 wireless link when different signal levels were present. They showed that without retransmissions implemented at the link layer, loss rates become unacceptable for any application [3]. The claim that MAC layer retransmissions improve TCP performance was confirmed also by *Xylomenos* and *Polyzos*, who experimented TCP and UDP on a WLAN and analyzed their behavior with different interfaces and bidirectional TCP traffic [4].

Said that, we may ask ourselves whether the current number of MAC layer retransmissions represents the optimal choice to support both TCP-based traffic and real time applications. Indeed, a high number of repeated retransmissions could still be not enough to prevent TCP from experiencing timeouts and retransmitting the same data as the MAC layer. At the same time, MAC retransmissions can be wasteful and potentially harmful for time-sensitive applications, such as real time video/audio or online games over UDP [166].

It should be said that a vast collection of research papers focusing on 802.11 could be found by delving into the Internet. They present analysis, problems, and solutions. Nonetheless, the vast majority of them provides results that focus on a throughput/losses point of view [57, 172, 173, 174, 176, 177], while the performance of real time applications depends on the measured per packet delay and jitter [58]. Even if some recent works present delay measurements for real time applications over IEEE 802.11, a deep analysis of this issue with respect to MMOG is still missing, as well as efficient solutions aimed at reducing queuing delay over wireless links [166, 175].

4 System Model

Aimed at providing a holistic solution for MMOGs and being aware of related works, we can now analytically study the four main requirements seen in Section 2.2 and the tradeoff relationship existing among them.

Every class of game is featured by a peculiar (and fixed) *Game Interactivity Threshold (GIT)* that represents the maximum delay endurable before visualizing a game event on players' screens if one wishes to preserve interactivity. The typical *GIT* for fast paced games (i.e. vehicle racing, first person shooter) corresponds to 150-200ms but this value can be increased up to seconds in case of slow paced games (i.e. strategic, role play game) [77, 78, 80, 106, 107].

If we call $t^{g(e)}$ the generation time of event e and $t_i^{v(e)}$ the visualization time of the same event at player i , then interactivity is preserved at i during the delivery of e when the following condition is satisfied:

$$t_i^{v(e)} - t^{g(e)} \leq GIT. \quad (1)$$

Both consistency and fairness regards having the same game state contemporary viewed in all the nodes of the system. Therefore, the same class of techniques is generally used to achieve each of them (or both). The easiest way to guarantee consistency and fairness is to make the game proceeding through discrete locksteps. Unfortunately, as discussed in Section 2.3, this scheme cannot be applied to interactive games.

To ensure fairness (and consistency) in continuously evolving games several studies propose schemes based on the introduction of artificial delays in order to contemporary visualize game events on all the players' screens (i.e., local lag schemes) [54, 78, 94, 104, 105].

With local lag, game advancements are delayed for a sufficient amount of time in order to guarantee that all the clients in the system process and perceive the generated game events at the same time and in the same order. Indeed, since the generation time of each event is unique and considering CC, the set of clients, we can say that we have event-related fairness [109] for event e if condition (2) is satisfied, simply stated, if there is a unique $t^{v(e)}$ value for all the players:

$$t_i^{v(e)} = t^{v(e)} \quad \forall i \in CC. \quad (2)$$

Since a single game event experiences different overall delays (OD) in its paths from the source to all the diverse destinations, different amounts of artificial delay δ should be added in order to contemporary visualize the same event e on all the players' screens and hence to satisfy the following condition:

$$t^{g(e)} + OD_i(e) + \delta_i(e) = t^{v(e)} \quad \forall i \in CC. \quad (3)$$

A possible value typically chosen for the unique $t^{v(e)}$ is represented by the highest OD in transmitting events amongst nodes. When the highest OD is greater than GIT , however, fairness is preserved at the cost of jeopardizing interactivity for all the players. Conversely, if we use GIT as an upper bound to $t^{v(e)}$, then we can guarantee interactivity but not fairness.

Consequently, in order to maximize the possibility to obtain both interactivity and fairness $t^{v(e)}$ should be set as

$$t^{v(e)} = t^{g(e)} + GIT. \quad (4)$$

The $OD_i(e)$ experienced by an event e when it finally reaches client i is composed by several delay components, respectively: physical latency $l_i(e)$, queuing time $q_i(e)$ on nodes along the path, and processing time $p_i(e)$. Therefore, $OD_i(e)$ can be written as

$$OD_i(e) = l_i(e) + q_i(e) + p_i(e). \quad (5)$$

Even when the network latency would allow having values of OD , and hence also of $t^{v(e)}$, lower than GIT , a large number of players generating a huge amount of traffic may raise the value of the other two components (i.e., $q_i(e)$ and $p_i(e)$), thus leading us again to the crossroad between fairness and interactivity.

To conclude, the efficiency and applicability of popular delayed-based algorithms such as local lag strongly depend on the network conditions and on the interactivity degree required by the game. Yet, guaranteeing both interactivity and full fairness through local lag can sometimes be achieved only at the cost of limiting the scalability of the game by bounding the number of contemporary participants and the geographical extension of the target player market.

It hence becomes evident how MMOGs require the use of architectural solutions and algorithms able to reduce the delay components in (5) in order to find the most efficient tradeoff among interactivity, consistency, fairness, and scalability.

5 Architectures

Typically, network architectures supporting MMOGs can be distinguished based on three main categories as depicted by Fig. 4: *centralized client-server*, *fully distributed*, and *mirrored game server*.

In the centralized client-server architecture, we have a single authoritative point which is responsible to run the main logic of the game. We report here only a few of its tasks: execute players' commands, enforce consistency, send back to the client the new game state update, etc. Clients have only to receive the new game state update, render it on the screen, and forward player's commands. The single authoritative point is usually represented by a single server; however, a cluster of computers could be utilized as well in order to increase the performance of the system [97].

The centralized client-server architecture represents the simplest solution for authentication procedures, security issues, and consistency maintenance [35, 95, 96, 97]. Moreover, assuming to have N simultaneous players, the generated messages are in the order of $O(N)$. On the other hand, the unique bottleneck limits the efficiency and scalability of this solution.

Fully distributed architectures are well represented by the peer-to-peer paradigm. In this case, all the involved nodes share the same intelligence and are responsible for running the whole logic of the system. In this case, in fact, each client has to autonomously update the game state view based on its player's commands and on game actions received from other players. This obviously requires terminals endowed with higher computational capabilities.

The main advantage in employing a fully distributed architecture is that of spreading the traffic load among many nodes thus generating a more scalable and failure resilient system [94, 112]. However, identical copies of the current game state need to be stored at each node. This requires some complex coordination scheme among peers; in fact, this scheme has to be distributed over the set of involved nodes and has to be able to guarantee the coherence of all game state views. Moreover, with fully distributed architecture, multicast should be employed to reduce the bandwidth requirements, but multicast technology is neither generally available nor mature enough for the specific application we are considering here. The exchanged messages could hence raise to the order of $O(N^2)$. Finally, authentication, cheating, and general consensus among all the peers are harder to be addressed than when a centralized architecture is employed.

Mirrored game server architectures represent a hybrid solution which efficiently embraces all the positive aspects of both centralized client-server and fully distributed architectures [59]. Based on this approach, GSSs are interconnected in a peer-to-peer fashion over the Internet and contain replicas of the same game state view. Players communicate with their closest GSS through the client-server paradigm. Each GSS gathers all the game events of its engaged players, updates the game state and regularly forwards it to all its players and GSS peers.

The presence of multiple high performance GSSs helps in distributing the traffic over the system and reduces the processing burden at each node [112]. Moreover, having each player connected to a close GSS reduces the impact of the player-dependent access technology (e.g., dial-up, cable, DSL) on the total delay experienced [110]. In this case, in fact, the communication among players results mainly deployed over links physically connecting GSSs, which can exploit the fastest available technology (e.g., optical fibers) to reduce latency. As a result, this architecture helps one in finding better solutions for the tradeoff among interactivity, consistency, fairness and scalability.

Other advantages in employing mirrored game server architecture are the absence of a single point of failure, the networking complexity maintained at server side, and the possibility to easily implement authentication procedures. Even if synchronization is still required to ensure the global consistency of the game state held by the various servers, this requirement is made easier with respect to fully distributed architectures thanks to the lower number of involved nodes. Assuming to have N players and M GSSs, for example, the generated game messages amount to $O(N+M)$, which is again $O(N)$ unless considering the unlikely case of having more servers than players.

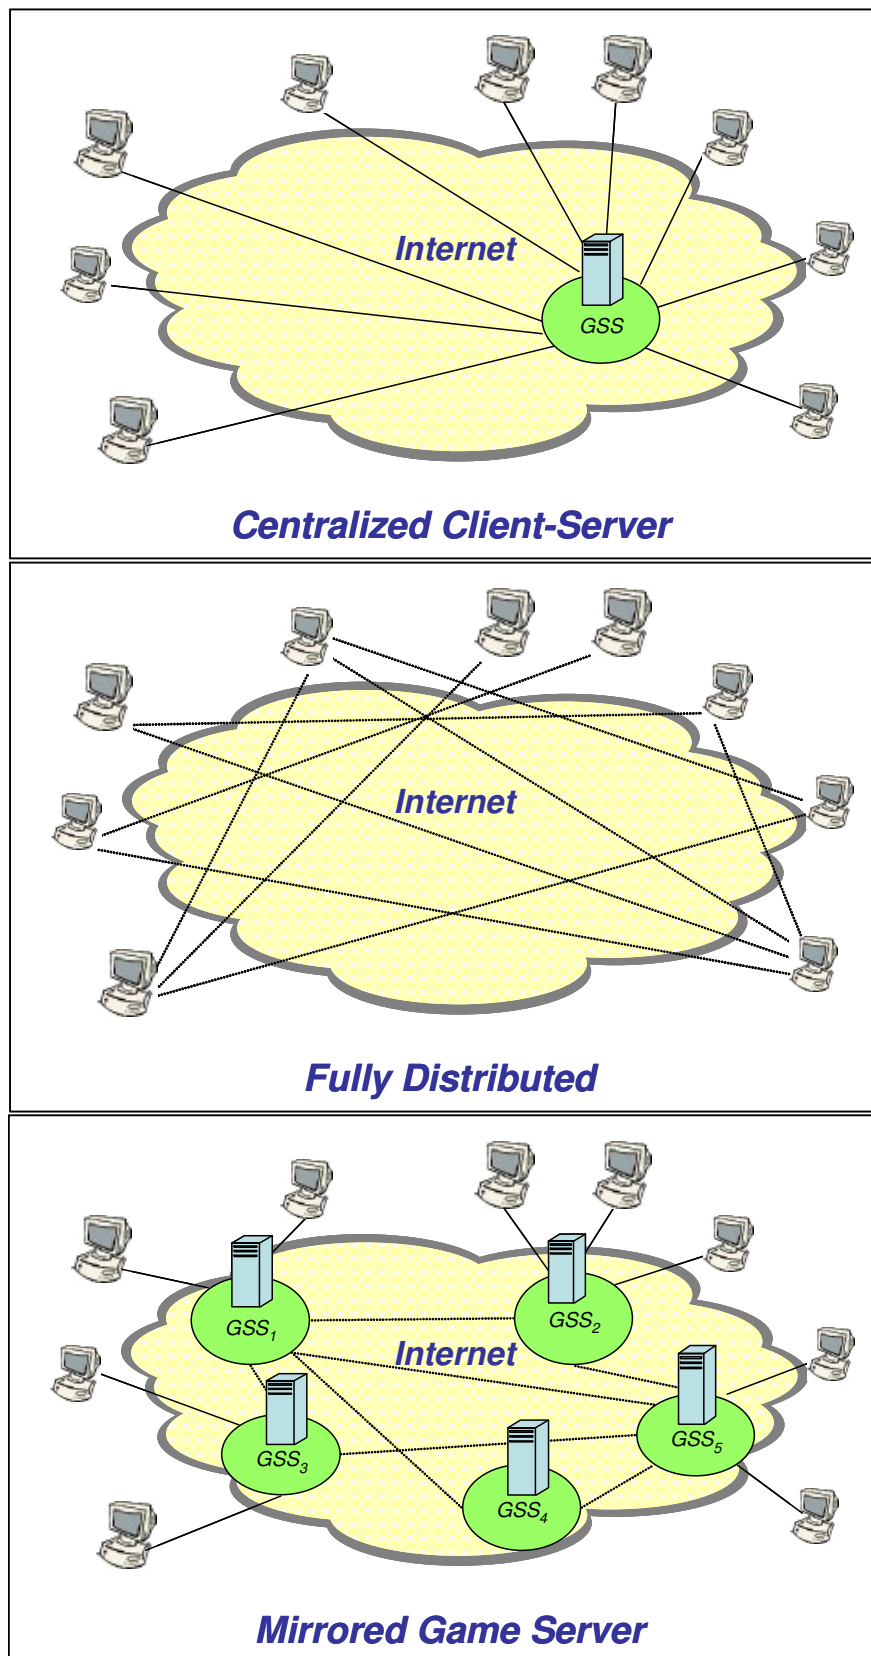


Figure 4. Online game architectures.

Chapter 3

Fast Synchronization Framework

1 Proposed Architecture

Since the analysis of advantages and disadvantages about the various possibilities illustrated in Section 2.5, mirrored game server emerges as the most appropriate architecture in order to efficiently manage large-scale distributed games. Indeed, this architecture embodies the advantages of both client-server and fully distributed paradigms and, in particular, it preserves the two most important features required by a MMOG architecture: scalability and controllability. The former is required to allow a multitude of players to engage in the same virtual arena. The latter regards the possibility to control the access to the game, avoid cheating, and have a centralized core where new techniques could be easily deployed to improve performance. Indeed, based on this architecture, we have devised an efficient synchronization scheme among GSSs able to enforce a high interactivity degree while guaranteeing a uniform view of the current game state.

Moreover, in Section 4 we will demonstrate how to exploit this architecture and our synchronization scheme to achieve also fairness, while in Section 5 we will focus on preserving interactivity and fairness also in the last hop of the connection. We will show how to enhance the mirrored game server architecture to make it able to sustain MMOG applications regardless of the connectivity type exploited on the client-server link.

2 Obsolescence and Correlation: Maintaining Responsiveness and Consistency

Absolute Consistency can be attained through the employment of a totally ordered event delivery scheme [60, 61]. On the other hand, this would imply an increment of the complexity and of the total delay experienced by the system. Waiting for the next in order action to be processed, while having other events ready in queue, may sensibly slow down the evolution of the game thus jeopardizing responsiveness.

Exploiting the semantics of the game can be put to good use to relax the total order delivery requirement and augment responsiveness [63, 64]. Some events, in fact, can lose their significance as time passes: new actions could make irrelevant previous ones. For example, player's movements are generally represented by final absolute position and, in case of rapid succession of movements of a single agent, the event representing its last destination makes obsolete the older ones.

Obsolescence can thus be defined as the relation between two received events e_1 and e_2 , generated at different times $t(e_1) < t(e_2)$, by which the existence of event e_2 diminishes the importance of processing also event e_1 , without affecting consistency (see Fig. 5-a). Dropping obsolete events before processing them clearly reduces computation at GSSs and speeds up the execution of fresher events.




Figure 5. Examples of obsolescence and correlation
 (a) Event e_2 makes obsolete e_1 ; (b) Event e_c is correlated to e_1 and rescinds the obsolescence of e_1 .

To define as obsolete a game event, we have to be sure that consistency would not be weakened. To this aim, we have also to take into account the notion of **correlation**. Two events, say e_1 and e_c , are correlated if the final game state depends on their execution order. Correlation has to be taken into account to determine the obsolescence of an event. As depicted in Fig. 5-b, it might be the case when e_2 would make obsolete a previous event e_1 but a further event e_c (correlated to e_1), temporary interleaved between e_1 and e_2 , rescinds this relationship of obsolescence. However, they are the only events that really need to be reliably delivered to all destined GSSs and in the same order as they were generated.

Total order delivery requirement can thus be relaxed in case of non-correlated game events. Their semantic independence, in fact, allows different GSSs to process them in diverse orders without affecting consistency. This means that non-correlated game events can be processed as soon as they are received without wasting any time in waiting preceding ones, thus augmenting responsiveness; missing game events can be interpolated and rendered on players' screens (see Section 2.3.1). Interested readers may find a deeper analysis on obsolescence and correlation in [63, 64].

On the other hand, our experiences with online games over a best effort network lead us to claim that there exist cases where even dropping all the obsolete events in a game is not enough to ensure an adequate interactivity degree. This is particularly true for a class of games that requires frenetic, and often redundant, players' actions. This class of games is widely recognized in the gaming community as fast-paced (or fast and furious) games: shoot/beat'em up games represent typical exemplars. For this class of games, even discarding some non-obsolete events may be preferred to compromise responsiveness and consequently users' perceived playability. Discarding some non obsolete event may generate some sporadic inconsistencies in the game state; nonetheless, we deem that partial consistency for a small amount of time becomes acceptable for the specific class of fast and furious online games. In

this scenario, in fact, the necessity of a very high interactivity degree emerges as overwhelming even on the full-consistency requirement.

3 Interactivity Restoring Mechanism

Exploiting the notions of obsolescence and correlation, Ferretti and Rocchetti propose a mechanism to restore interactivity in mirrored game server architecture [63]. The player's actions are collected by the closer GSS, transformed into events and finally forwarded to other GSSs in order to maintain a global identical view of the game state. Events are marked at their creation with a generation timestamp and then sent to destination: they are hence orderable. Obviously, a global concept of time has to be maintained by all the GSSs. Different solutions proposed in literature can synchronize the clocks at each GSS [65, 66, 67]. Alternatively, this can also be obtained by exploiting some new technological synchronization device such as GPS.

Since UDP is used as the transmission protocol, missing packets are handled at the application level. Game state updates transmitted by GSSs, in fact, are identified by sequence number. By the means of NACK (Negative ACKnowledgment) packets it is possible to determine which ones are missing and decide whether to retransmit them or not.

Each receiving GSS considers the arrival time of the event and measures the difference elapsed since its generation; the resulting value is named *Game Time Difference (GTD)*. The *GTD* of the event is then compared with the predefined constant *GIT* (see Section 4) and normal delivery operations are performed until the former value is lower than the latter. When the *GTD* value exceeds the *GIT*, the GSS turns on a stabilization mechanism which exploits the obsolescence notion to drop useless events so as to bring the *GTD* back within the *GIT*. Moreover, messages are sent to the connected GSSs in order to make them aware of the witnessed lack of responsiveness. These GSSs can thus avoid to forward obsolete and missing events to the non interactive GSS thus helping in restoring faster

4 RED/RIO Techniques

Random Early Detection (RED) algorithm is an active congestion avoidance mechanism enforced at routers [68]. Traditional queue management employs simple "tail drop" schemes that drop packets only when the queue overflows. Conversely, RED algorithm randomly performs early packet discards to notify sources about the incipient congestion. In this way, a single loss experienced by a sender smoothly decreases the entire congestion level of the network and keeps the average queue size at a low level. The rationale lies in the gained capability of better accommodating occasional bursts of packets and avoiding situations in which several connections simultaneously decrease their sending rate. Summarizing, RED avoids severe congestion and maintains a stable traffic level in place of dealing with congestion after already occurred.

Every time the router receives a packet, the RED algorithm calculates the new average queue size and the probability to discard the packet. The computing method utilizes a uniform random variable, which has been proven to be more adequate to this aim than a geometric random variable. In fact, a uniformly distributed discarding function avoids global synchronization thus attaining an unwavering course of transmissions. The dropping probability is bounded by two thresholds of the queue size: min and max. Within this interval, the probability to drop a packet increases from 0 to a maximum discarding probability (P_{max}). Under min, no packet is dropped and beyond max all packets are discarded.

RIO (RED with In and Out) scheme is an enhanced version of RED mechanism able to discriminate between two different classes of traffic, non-prioritized (Out) and prioritized (In), and calculates two distinct dropping probabilities [70]. The two dropping functions work independently and are featured with specific boundaries and slopes. They hence permit to discard packets utilizing different probabilities and congestion levels depending on the packets' traffic class.

5 Enhancing Interactivity with RED/RIO Techniques

Taking inspiration from the RED approach in case of incipient congestion in best effort networks, we have recently enhanced the aforementioned Interactivity Restoring mechanism with the Interactivity-Loss Avoidance (ILA) approach [62]. The main innovation is the capability to preempt responsiveness disruption instead of restoring it after having already lost playability. To this aim, the system discards some packets when the responsiveness among GSSs descends significantly. In practice, ILA replaces the basic binary dropping mechanism for obsolete events (OFF when interactivity is present and ON when interactivity is lost) with a continuously-working proactive mechanism that drops obsolete events with a probability that depends on the level of responsiveness of the system.

Even if, similarly to RED, ILA utilizes a uniformly distributed dropping function, the parameter taken under control is the average *GTD* instead of the average queue size. Upon each packet arrival, in fact, each GSS determines the *GTD* of the arrived event, namely *sample_GTD*, and feeds a low pass filter to compute the updated average *GTD*, namely *avg_GTD*. When *avg_GTD* exceeds a certain threshold, the GSS drops obsolete events with a certain probability p , without processing them. If *avg_GTD* exceeds a subsequent limit, p is set equal to 1, and all obsolete events waiting for being processed are discarded.

Obviously, if the low interactivity degree perceived by some GSSs is not affecting other GSSs, the latter can process and forward to their clients even those packets discarded by the formers. Indeed, since obsolete events are actions that can be considered non-critical for the game evolution, there is no unfairness generated by processing them only in some GSSs. On the other hand, not processing those events in highly interactive GSSs would represent an unjustified limitation, even if light, to the number of events visualized at some player's side. For instance, an intermediate movement of a game character, when the final position is further ahead and no shooting nor other hazards could harm it, could receive different treatments: it can be dropped by some GSSs to preserve interactivity, or can as well be processed by other GSSs for the sake of a complete and fluent game visualization on the screens of their clients without relying on imprecise and computationally expensive interpolations.

To ensure an adequate playability degree even to the class of fast and furious games we have then further enhanced our ILA scheme. For this class of games, in fact, the core attractiveness for players emerges from a feverish, sometimes even chaotic, action sequence of user's actions. We have hence enhanced our ILA scheme with features derived from the integration of a RIO-like algorithm in place of the RED-like one [69]. The additional dropping probability provides the possibility to discard even non-obsolete game events when dropping all the obsolete ones is not yet sufficient to maintain an adequate level of responsiveness. The two discarding functions are featured with specific parameters; they work independently one from the other and take action in sequence with the increasing of the game event *GTDs* at the GSSs.

Dropping non-obsolete events can be done without consequences only for a category of games where little inconsistencies are not highly deleterious for the aim of the game and for player's fun (e.g., fast-paced games). Even in this case, if the number of dropped non-obsolete events becomes significant, a consistency restoring mechanism may be required to re-establish a coherent game state view among all GSSs [71]. For the sake of clarity, from here on we are going to call ILA-RED the ILA version that discards only obsolete events, while ILA-RIO represents the version with two discarding functions. When we want to refer to both the algorithm with no distinction we just use the word ILA.

In Fig. 6 we depict the two discarding functions of ILA-RIO. Three parameters (and three phases) characterize each of the twin algorithms: \min_o , \max_o and $P\max_o$, for obsolete events, and \min_v , \max_v and $P\max_v$ for valid (i.e., non obsolete) ones. In the graph, the y-axis represents the dropping probability corresponding to the *avg_GTD* indicated by the x-axis. Focusing on obsolete events, for values of *avg_GTD* in $[0, \min_o)$ the mechanism performs normal operations, with no packet drops, while in $[\min_o, \max_o)$ obsolete packets are discarded with a computed probability, and finally in $[\max_o, \infty)$ all obsolete packets are thrown away. The intervals $[0, \min_v)$, $[\min_v, \max_v)$, and $[\max_v, \infty)$ define the corresponding phases for valid events. The dropping probabilities are computed as a function of *avg_GTD* and of $P\max_o$ or $P\max_v$, respectively. Persistent situations of low interactivity degree result in high *avg_GTD* and hence in high discarding probabilities. High dropping probability values (for $P\max_o$ or $P\max_v$) will make the GSS discarding events without processing nor forwarding them, thus helping in restoring an adequate level of time interaction between servers.

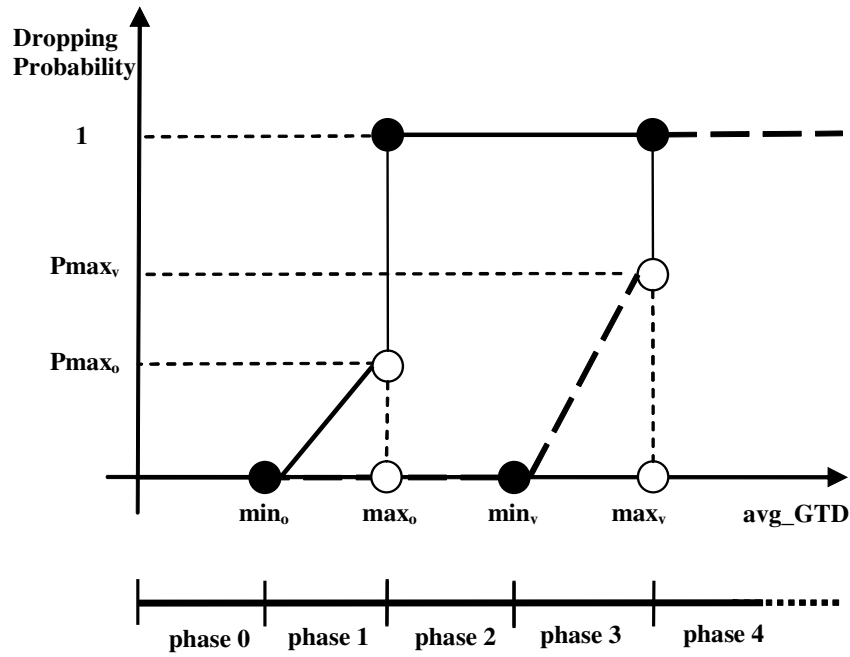


Figure 6. Discarding probability functions for ILA-RIO.

```

0] for each event_packet arrival {
1]   determine the sample_GTD
2]   calculate the new average delay avg_GTD
3]   if ( $min_o \leq avg\_GTD < max_o$ ) then
4]     calculate the probability  $P_o$  of dropping an obsolete event
5]     determine if ONE obsolete event has to be discarded
6]   else if ( $max_o \leq avg\_GTD$ ) then
7]     drop ALL obsolete events
8]   if ( $min_v \leq avg\_GTD < max_v$ ) then
9]     calculate the probability  $P_v$  of dropping a valid event
10]    determine if ONE valid event has to be discarded
11]   else if ( $max_v \leq avg\_GTD$ ) then
12]     drop ALL valid events
13]   endif
14] endif
15] endfor

```

Figure 7. ILA-RIO algorithm.

Since valid events are strictly linked to consistency, the possibility to discard them should be taken into account only as last resort, in case of heavy disruption of responsiveness. For this reason, ILA-RIO starts dropping obsolete packets much earlier than valid ones. In addition, the algorithm throws away all the obsolete packets before considering any dropping probability on valid events; simply stated: \max_o is smaller than \min_v . Finally, diverse aggressiveness in dropping packets, depending on their class, can be decided by adjusting the values of P_{\max_o} and P_{\max_v} .

In essence, the algorithm repeats a block of operations (listed in Fig. 7) each time a new event arrives at the considered GSS. In particular, the *GTD* of the game event is calculated (*sample_GTD*, line 1) as the difference between the generation time at the sending GSS and its delivery time to the considered receiving GSS. The scheme feeds a low pass filter with the just calculated *sample_GTD* in order to update the average of the *GTDs* (*avg_GTD*, line 2):

$$\text{avg_GTD} = \text{avg_GTD} + w * (\text{sample_GTD} - \text{avg_GTD}). \quad (6)$$

In (6), w is a sensitivity coefficient, with values comprised in $(0, 1]$, that determines how closely the trajectory of the average follows the movements of the samples. Higher values of w correspond to heavier relative weights of the last sample in the current average.

While *avg_GTD* lies below \min_o , the process stays in phase 0 and no particular operation is performed. Conversely, when *avg_GTD* is comprised between \min_o and \max_o , then the scheme is in phase 1 and lines 4-5 are executed. Basically, a dropping probability P_o is computed as a fraction of P_{\max_o} in order to establish whether an obsolete event must be discarded; this fraction linearly corresponds to the position of *avg_GTD* in the interval $[\min_o, \max_o)$.

$$P_o = \frac{P_{\max_o} \times (\text{avg_GTD} - \min_o)}{(\max_o - \min_o)}. \quad (7)$$

Analogously, when *avg_GTD* is comprised between \min_v and \max_v , then the scheme lies in phase 3 and lines 9-10 are executed. The probability P_v to discard a valid event is then calculated as a fraction of P_{\max_v} depending on the position of *avg_GTD* in the interval $[\min_v, \max_v)$.

$$P_v = \frac{P_{\max_v} \times (\text{avg_GTD} - \min_v)}{(\max_v - \min_v)}. \quad (8)$$

If the algorithm stays in phase 1 or 3 the corresponding dropping probability increases until an event is discarded. In particular, two couple of constants, $L1$ and $L2$ for P_o , $H1$ and $H2$ for P_v , are defined as shown in (9), (10), (11), and (12), respectively.

$$L1 = \frac{P_{\max_o}}{(\max_o - \min_o)}, \quad (9)$$

$$L2 = \frac{P_{\max_o} \times \min_o}{(\max_o - \min_o)}, \quad (10)$$

$$H1 = \frac{P_{\max_v}}{(\max_v - \min_v)}, \quad (11)$$

$$H2 = \frac{P_{\max_v} \times \min_v}{(\max_v - \min_v)}. \quad (12)$$

Utilizing these constants to rewrite (7) and (8), we obtain (13) and (14) which are employed respectively in line (4) and (9) of the ILA-RIO algorithm.

$$P_o = L1 \times \text{avg_GTD} - L2 \quad (13)$$

$$P_v = H1 \times \text{avg_GTD} - H2 \quad (14)$$

So far, the two probability functions result in geometric random variable distributions of the drops, while it would be desirable to discard events at fairly regular intervals. This can be obtained by applying an incrementing counter to augment, at each iteration, the weight of the considered discarding probability; as result, we generate a uniform distribution of the dropping dispersion, which makes the ILA scheme more resilient to temporary bursty periods [68]. To this aim, when the algorithm is in phase 1 or 3, P_o or P_v is compared with a number which is randomly generated within the interval $[0, 1)$. If we call R_o and R_v , the two random numbers, we can see from (15) and (16) how they are used to take into account the number of iterations elapsed since the last drop to increment the discarding probability. Basically, in phase 1 an obsolete event is dropped (line 5) if

$$counter_o \geq (R_o / P_o). \tag{15}$$

Analogously, in phase 3 a valid event is dropped (line 10) if

$$counter_v \geq (R_v / P_v). \tag{16}$$

The variables $counter_o$ and R_o are reinitialized to respectively 0 and a new random value each time the algorithm enters phase 1, or some obsolete packet is discarded. Analogously, variables $counter_v$, and R_v are reinitialized each time the algorithm drops a valid packet. Conversely, if the scheme is respectively in phase 1 or 3, the two dropping probabilities, P_o and P_v , are recomputed every new event arrival. Moreover, $counter_o$ and $counter_v$ are incremented by 1 every time ILA-RIO stays in phase 1 or phase 3, respectively, without dropping any game event.

Considering the remaining of the algorithm, if avg_GTD grows beyond max_o , the scheme enters in phase 2 and all obsolete packets have to be discarded in the attempt of re-establishing interactivity (line 7). Moreover if the value of avg_GTD surpasses even max_v , all the events, with no distinction between obsolete or valid, are dropped (line 12). At this point, in fact, the extremely jeopardized responsiveness conditions suggest resetting and reinitializing the game session.

Finally, it is easy to notice that ILA-RED can be considered a particular instance of ILA-RIO. In fact, if we set to 0 the probability of dropping a non obsolete event, or to ∞ the GTD value at which it takes action, then the two schemes behave identically. Simply stated, a chart representing ILA-RED behavior, equivalent to that one presented in Fig. 6, would contain just phase 0, phase 1 and phase 2, respectively contained in the intervals $[0, min_o)$, $[min_o, max_o)$ and $[max_o, \infty)$. Analogously, the algorithm for ILA-RED can be obtained by simply eliminating lines 8-13 from Fig. 7.

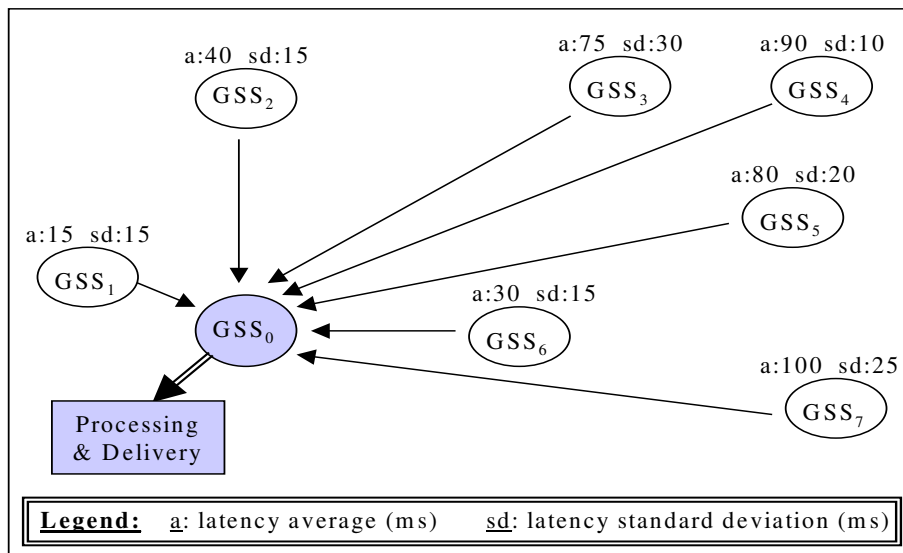


Figure 8. The adopted configuration.

TABLE I. SENDING GSSs INVOLVED IN THE SIMULATIONS.

Number of Sending GSSs	Corresponding GSSs employed
4	GSS ₁ , GSS ₂ , GSS ₃ and GSS ₄
5	GSS ₁ , GSS ₂ , GSS ₃ , GSS ₄ and GSS ₅
6	GSS ₁ , GSS ₂ , GSS ₃ , GSS ₄ , GSS ₅ and GSS ₆
7	GSS ₁ , GSS ₂ , GSS ₃ , GSS ₄ , GSS ₅ , GSS ₆ and GSS ₇

6 Simulation Assessment

To evaluate our event processing strategy, we have created a mathematical model simulating a general Mirrored Game Server architecture comprising various GSSs connected via diverse links over the Internet. As said, we assume that the events generated in the system can be totally ordered by exploiting a global notion of time.

Different numbers of GSSs have been simulated in our scenario. In order to permit a deeper analysis of the dynamics related to game event traffic and processing, simulations have focused on the packet receiving aspect of a single GSS, while the other GSSs are involved as game event transmitters. Fig. 8 depicts the adopted configuration of the network and shows the values assigned to the simulation parameters. GSS₀ is the receiving GSS and the others are the sending GSSs. We carried out several simulation experiments (30) with a number of sending servers varying in the range from four to seven. The involved GSSs, for each different configuration, are listed in Table I.

We have been inspired by the scientific literature on MMOGs to set the values of the network latencies among the GSSs [72]. Specifically, GTDs follow a lognormal distribution having the average and standard deviation values as shown in Fig. 8. If each client connects to the “nearest” GSS to bound the impact of the client-server latency on the total delay experienced by the game events, we can assume to have the client-server latency portion already comprised in the network latency values mentioned above. Finally, the event generation rate at each GSS and the average event size (200 Bytes) have been chosen as inspired by the games literature and varies from a normal traffic situation to an intensely loaded one [73].

We define as Average Inter-Departure Time (AIDT) the time that elapses, in average, between two subsequent game packet departures at a single server. The AIDT values have been used to generate the lognormal distribution of the event departures from each involved server. Various simulations have been run having the AIDT parameter set to 30ms, 45ms or 60ms. The standard deviation was always equal to 10ms. The chosen AIDT values produce an amount of events as those typically generated by from 5 to tens of players for each given GSS. Based on these values, we have generated three diverse trace files containing 1000 events for each GSS. Each trace file also included the information needed to identify (correlated and) obsolete events. Where not differently stated, we have set to 90% the probability that an event makes obsolete preceding ones. This represents a realistic scenario for a vast plethora of possible games where critical (correlated) game events that cannot become obsolete have to be considered only sporadically, such as during collisions or shots, and may represent even less than the 10% of the whole set of game events.

As a confirmation of this claim, an extensive study of players’ behavior on Quake 3 is presented in [76]. In that paper, a measure of the average number of kill actions per minute as a function of the median ping time between client and server is reported. Using those measures we can provide a numerical explanation that demonstrates how 10% of correlated events and 90% of obsolescence probability may represent a realistic scenario for interactive MMOGs [81].

Specifically, in [76] it is shown that when the 80% of players is located within 180ms of range from the server, the average of kill actions per minute varies in the interval from 1.60 to 3.25. Considering the median of this interval (2.425) we obtain 0.04 kill events per second. With an AIDT of 60ms at client

side [73, 77], 16.67 game events are generated every second. Therefore, the resulting percentage of kill events over the whole set of game actions amounts to just 0.24%. Pessimistically assuming that a kill event can be issued only after an average number of 40 correlated actions (e.g., various shots, movements of the character into the location where it will be shot or out of the position where it would have been shot) we get 9.6% of critical events. Therefore, 10% of correlated events and 90% of obsolescence probability represent a realistic configuration for online games simulations.

Since real commercial games employ UDP as the transport protocol, we have adopted it too for our simulations [75]. Moreover, to circumvent the problems deriving from UDP's unreliability, we have implemented an application level retransmission scheme based on NACKs (Negative ACKnowledgments).

Each experiment has been replicated to compare the outcomes of three different synchronization schemes: our proposed ILA-RED scheme, the ON-OFF mechanism (Interactivity Restoring as reviewed in Section 3.3), the traditional OFF approach (having no discrimination of obsolete packets and no event discarding nor other algorithms to restore interactivity). We have then also run some preliminary experiments involving ILA-RIO.

As to the algorithm, we have chosen to set $w=1/8$ in (6) in the attempt to make the algorithm able to filter out sporadic high *GTD*s, while being able to promptly react to a persistent decline of responsiveness. In the ILA-RED case, we have set parameters as follows: $\min_o = 50\text{ms}$, $\max_o = 150\text{ms}$ (equivalent to the *GIT* for the ON-OFF scheme) and $P_{\max_o} = 0.2$. Parameters in ILA-RIO were chosen as $\min_o = 50\text{ms}$, $\max_o = 100\text{ms}$, $P_{\max_o} = 0.2$, $\min_v = 150\text{ms}$ (equivalent to the *GIT* for the ON-OFF scheme), $\max_v = 225\text{ms}$ and $P_{\max_v} = 0.3$.

Choosing appropriate values for the above cited parameters is an important part of implementing ILA algorithm. The various phases, in fact, should take place in precise situations experienced by the game platform. In particular, phase 1 have to be activated when the delay between the generation of a player's action and its execution on the screens provides the first perceivable symptoms of responsiveness degradation. Since phase 2 corresponds to dropping all the obsolete events, it should take action when the lag becomes annoying and low-performance determining for players. In case of fast paced games, the parameter of these two phases should be chosen in a further conservative way, in order to anticipate the triggering times for the discarding functions. Phase 3 represents the last resort to restore responsiveness in the system, whilst still having latencies that can be temporary tolerated by players. Finally, since entering in phase 4 amounts to reinitialize the game session, its entering threshold have to be chosen as the higher latency value at which users still can compensate the poor playing conditions by anticipating their moves.

The various thresholds should be hence set differently depending on the played game or at list on the belonging game class (e.g. adventures, shoot/beat'em up, car racing, etc.). However, as rationale for our chosen values, scientific literature declares that a delay of 50ms is not perceived at all by players while at 150ms (i.e., our *GIT* parameters) player's performance results disturbed by the lag. Finally, 225ms of delay could represent an upper bound for playable interaction [76, 77, 78, 79]. These limits hold for games like vehicle racing, first person shooters and fast shoot/beat'em-up, but can be augmented in case of strategic games (e.g. Starcraft, Age of Empire, etc.) [80].

The interactivity benefits attained by obsolete event discarding schemes can be analyzed in terms of delays experienced by game events before being processed. Therefore, we have compared the responsiveness of all the three schemes considered (i.e. ILA-RED, ON-OFF, OFF), exploiting the following metrics: i) the number of events having a *GTD* larger than the *GIT*; ii) the cumulative function of the *GTD* values; iii) the average of the *GTD* values, their standard deviation, the minimum and maximum values.

The fluency of the game evolution on players' screens passes through having just a very limited number of discarded game events. We have hence compared the two obsolete event discarding schemes (i.e. ILA-RED, ON-OFF), measuring: i) the number of obsolete events dropped; ii) the number of times all the pending obsolete events have to be discarded; iii) the average number of subsequent game events with *GTD* over *GIT*; iv) the total number of bursts of game events with *GTD* over *GIT*.

In order to deeply understand the benefits attainable by employing a proactive mechanism as ILA-RED, under different levels of stress of the game platform, we have performed a sensitivity analysis depending on the considered AIDT parameter. In particular, we have compared ILA-RED and ON-OFF, under different values of AIDT, by measuring: i) the number of obsolete events dropped; ii) the average of the *GTD* values, their standard deviation, and the minimum and maximum values.

Finally, some preliminary results about ILA-RIO have been obtained by comparing it against ON-OFF and OFF mechanisms under different conditions. Specifically, considering different probabilities of obsolescence among game events, we have measured: i) the number of events having a GTD larger than the GIT ; ii) the number of events dropped, both obsolete and valid ones.

7 Current Results

We present here the most relevant results among those we obtained.

In particular we compare the different schemes by showing the outcomes of simulations having the AIDT parameter set to 30ms, 45ms, and 60ms, respectively. Having a low AIDT corresponds to a higher frequency in game event transmissions at each server and therefore in a higher congestion and computational workload.

7.1 Obsolescence-Based Scheme vs Traditional One

As predicted, the two obsolescence-based discarding mechanisms outperform the traditional one, especially with high traffic load at the servers. In fact, Fig. 9 compares, for ILA-RED, ON-OFF and OFF schemes, the percentage of events arrived at GSS_0 with a GTD value larger than the GIT . This also represents the percentage of events that cannot be considered interactive.

The cumulative function of the $GTDs$ (Fig. 10) represents another tool proficient in evaluating the efficacy of ILA-RED and ON-OFF schemes. Indeed, the more the line is concentrated in the left side of the chart, the higher is the percentage of events having a GTD lower than a certain threshold. In particular, Fig. 10 depicts the cumulative function of the $GTDs$ in a scenario considering seven sending $GSSs$, each one sending events to the receiving GSS_0 with an AIDT of 30ms. In this configuration, ILA-RED has 93.86% of events with a GTD less or equal than the GIT of 150ms, ON-OFF hits the 89.40%, and OFF reaches only the 49.94%.

These results are coherent with the values of the average and the standard deviation of the $GTDs$ considering all the events transmitted. Table II shows sensible reductions in the values of both these metrics when ILA-RED or ON-OFF are implemented. Moreover, the two obsolescence-based discarding schemes result more resilient to an increased event generation activity within our mirrored game server architecture. This is evident if the case of seven sending $GSSs$ is compared with the one employing only four sending $GSSs$. In this case, the average of the $GTDs$ decreases from 19.72% (OFF) to 12.07% (ON-OFF) and 11.71% (ILA-RED).

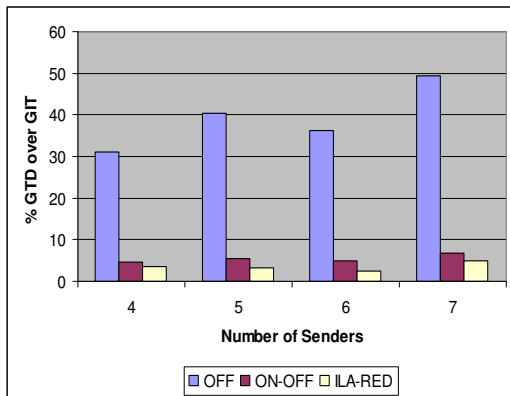


Figure 9. Percentage of events with GTD over GIT; AIDT = 30ms.

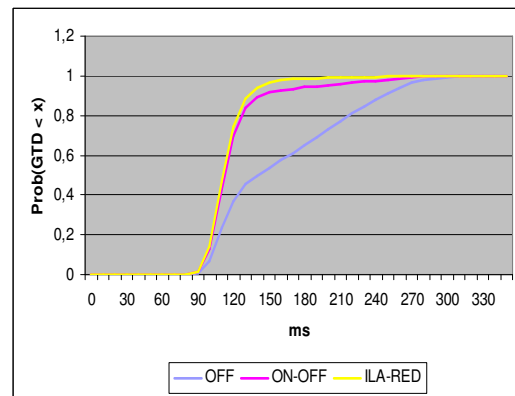


Figure 10. Cumulative function of the GTDs in a scenario with 7 GSSs; AIDT = 30ms.

TABLE II. MAXIMUM, MINIMUM, AVERAGE AND STANDARD DEVIATION OF THE GTDs (MS); AIDT = 30MS.

	4 GSSs			5 GSSs			6 GSSs			7 GSSs		
	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED
MAX	324	324	325	325	324	277	318	319	278	345	345	300
MIN	88	88	86	88	88	88	87	88	88	93	93	93
AVG	142	116	111	153	120	115	148	119	114	170	130	124
ST.DEV	52	30	20	53	32	19	50	28	18	56	32	19

7.2 ILA-RED vs ON-OFF: a Comparative Evaluation

The proactive mechanism implemented by ILA-RED does not just slightly improve the already good performance achieved by the ON-OFF scheme. Indeed, Table II also shows that the standard deviation of the *GTDs* attained employing ILA-RED is always sensibly smaller than the one obtained utilizing ON-OFF. The game actions flow more homogeneously on the screen thus providing a more pleasant game experience for the user.

Not only obtains ILA-RED a slightly better interaction level with respect to ON-OFF, but the total number of discarded events to attain this positive result is definitively lower. Again, with the AIDT parameter equal to 30ms at each GSS, we can analyze the behavior of the compared synchronization schemes in the challenging situation of a stressed game platform. To this aim, Fig. 11 shows that the results in Fig. 9 and Table II are obtained by ILA-RED at the cost of circa only 40% of the obsolete events dropped by ON-OFF. In essence, by anticipating some event drops, ILA-RED preempts the loss of an acceptable interactivity degree and smoothes the experienced *GTDs*.

Discarding too many obsolete events can result in sudden jumps and temporary interruptions of the images/video flow on the player's screen. Even interpolation experiences an increasing number of errors when gaps are too wide and the consequent jerky rendering could be very annoying for customers. Drops should hence be avoided whenever possible. Since ILA-RED needs less pervasive interventions to be effective, we can say that even if both schemes ensure responsiveness and consistency, ILA-RED outperforms ON-OFF and finds an efficient tradeoff between the percentage of obsolete events to be discarded and a fluent visual progression of the game.




Figure 11. Percentage of discarded events; AIDT = 30ms.




Figure 12. # of activations of phase ON and phase 2 for ON-OFF and ILA respectively; AIDT = 30ms.




Figure 13. Average size of non-interactive bursts; AIDT = 30ms.

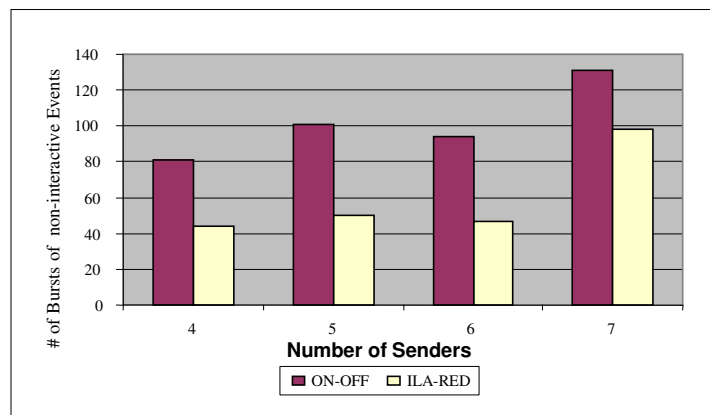


Figure 14. Total Number of bursts of non-interactive events; AIDT = 30ms.

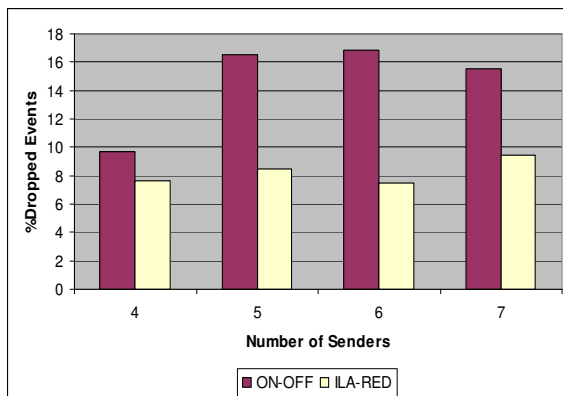


Figure 15. Percentage of discarded events; AIDT = 45ms.




Figure 16. Percentage of discarded events; AIDT = 60ms.

The positive effects gained by the probabilistic preemptive discarding mechanism are further highlighted by Fig. 12. The columns in the picture correspond to the number of times the employed scheme resorts to dropping obsolete events to restore the disrupted responsiveness. As it is evident from the great difference between the columns, the preventive probabilistic drop of some game events in phase 1 strongly reduces the number of times ILA-RED needs to discard all the obsolete packets.

Another important metric to better evaluate synchronization schemes is concerned with the distribution of the event drops throughout the game. Fig. 13 and Fig. 14 contribute to provide further insight along this direction. In particular, Fig. 13 shows the average number of consecutive game events with an instantaneous *GTD* value larger than *GIT* (average size of non-interactive bursts). Fig. 14, instead, shows the total number of non-interactive bursts obtained during our simulations. As it is evident from these two charts, while ILA-RED and ON-OFF schemes have an almost equal average size of bursts of non-interactive events, the former allows only a reduced number of these bursts, thus confirming the benefits announced by Fig. 11.

7.3 ILA-RED: Sensitivity Analysis

It is evident from the previous sections that ILA-RED results more efficient than ON-OFF scheme in ensuring responsiveness whilst minimizing the number of drops. We deem that this advantage becomes more evident when the game event traffic is more intense, while the two synchronization mechanisms behave the same when plunged into a low stress environment. Simply stated, our hypothesis is that the larger the AIDT value, the more ILA-RED and ON-OFF become equivalent in terms of the number of dropped events required to maintain a proper interactivity degree. We intend to find the breakeven point in the performance curves between ILA-RED and ON-OFF. For this purpose, we have run additional simulations with the lognormal distribution of the departing game packets from each server generated with the AIDT value equal to: respectively, 45ms and 60ms.

With this greater AIDT value both the links and the GSSs result less congested. As a result, the difference in the number of discarded game events between ILA-RED and ON-OFF progressively diminishes (see Fig. 15 and Fig. 16). With AIDT = 60ms, the number of discarded events is almost the same for the two synchronization schemes. At this point, however, also the large advantage against the OFF scheme has been sensibly reduced (see Fig. 17 and Fig. 18); despite of this, better average and standard deviation of the *GTDs* still persist (see Table III and Table IV).

This phenomenon has two clear explanations. First, having lower game event generation rates decreases both the latencies in the network and the queuing time at the receiving GSS, thus naturally improving the interactivity degree yet without any external intervention. Second, we have to remind that both ILA-RED and ON-OFF schemes rely on discarding obsolete events when the interactivity level decreases, without wasting time in processing them. Consequently, to be effective, these schemes require to have droppable obsolete events queued at the receiving GSS while the server is impacted by processing delays; this engulfment of events waiting for being processed is more likely to happen with lower AIDT values.

TABLE III. MAXIMUM, MINIMUM, AVERAGE AND STANDARD DEVIATION OF THE *GTDs* (MS); AIDT = 45MS.

	4 GSSs			5 GSSs			6 GSSs			7 GSSs		
	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED
MAX	331	331	331	331	332	331	315	315	281	338	331	332
MIN	87	87	87	87	87	87	89	89	89	91	91	91
AVG	139	120	120	148	128	125	145	128	125	164	139	135
ST.DEV	45	24	22	46	27	22	43	27	20	49	27	21

TABLE IV. MAXIMUM, MINIMUM, AVERAGE AND STANDARD DEVIATION OF THE GTDs (MS); AIDT = 60MS.

	4 GSSs			5 GSSs			6 GSSs			7 GSSs		
	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED	OFF	ON-OFF	ILA-RED
MAX	328	326	326	328	325	325	319	286	286	340	328	325
MIN	88	88	88	88	88	88	87	87	86	92	91	90
AVG	144	133	133	154	142	142	153	142	142	169	155	154
ST.DEV	42	29	29	41	29	28	40	29	29	42	27	27

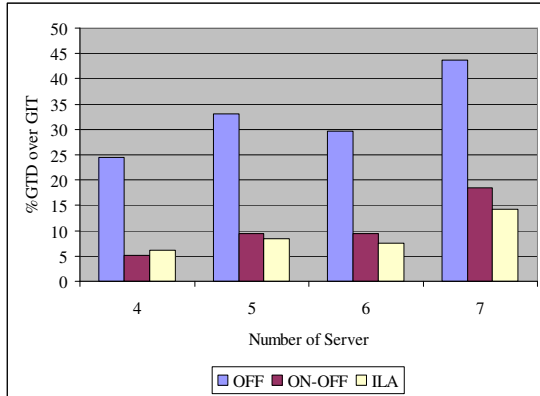


Figure 17. Percentage of events with GTD over GIT; AIDT = 45ms

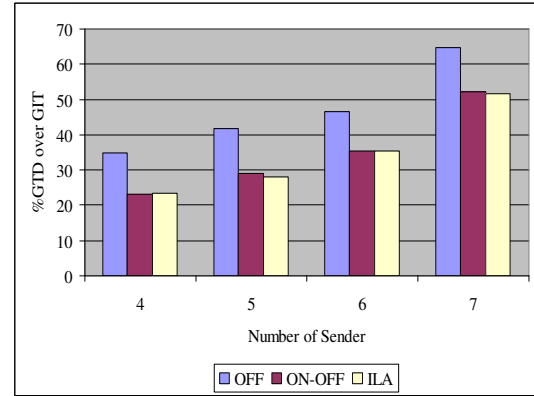


Figure 18. Percentage of events with GTD over GIT; AIDT = 60ms.

7.4 ILA-RIO: Some Preliminary Results

We present some preliminary results related to the use of the ILA-RIO scheme to ensure high responsiveness even with fast paced games. In particular, Fig. 19 and Fig. 20 refer to specific event trace configurations, respectively utilizing different probabilities of obsolescence among events.

As observable (Fig. 19-a, Fig. 20-a), in both configurations the obsolescence-based schemes outperform again the traditional OFF method in terms of *GTDs*. Moreover, ILA-RIO further diminishes the number of events that are featured with a *GTD* larger than the predetermined *GIT* w.r.t. ON-OFF. The game evolution fluency results improved as well since ILA-RIO greatly reduces the amount of dropped events required (Fig. 20-b). This difference disappears in Fig. 19-b; having a smaller probability of obsolescence (50%), in fact, reduces the number of available obsolescent packet that could, in case, be discarded. Thereby, both for ILA-RIO and for ON-OFF, even discarding all the obsolete queued events could still not be sufficient to promptly restore the required high interactivity degree with the considered intense traffic load. Indeed, this is exactly the case when dropping even valid events may become acceptable for fast paced games.

Finally, we evaluated the amount of valid game events dropped by our ILA-RIO approach. Table V reports the percentage of obsolete and valid events that are discarded, depending on the event trace. As expected, the amount of dropped valid game events diminishes as the percentage of obsolete ones becomes greater. This tendency is due to the fact that if an adequate number of obsolete events are available during the events exchange activity, then our scheme can exploit the drop of all these (obsolete) events to restore responsiveness when entering phase 2, with no need of resorting to valid ones.

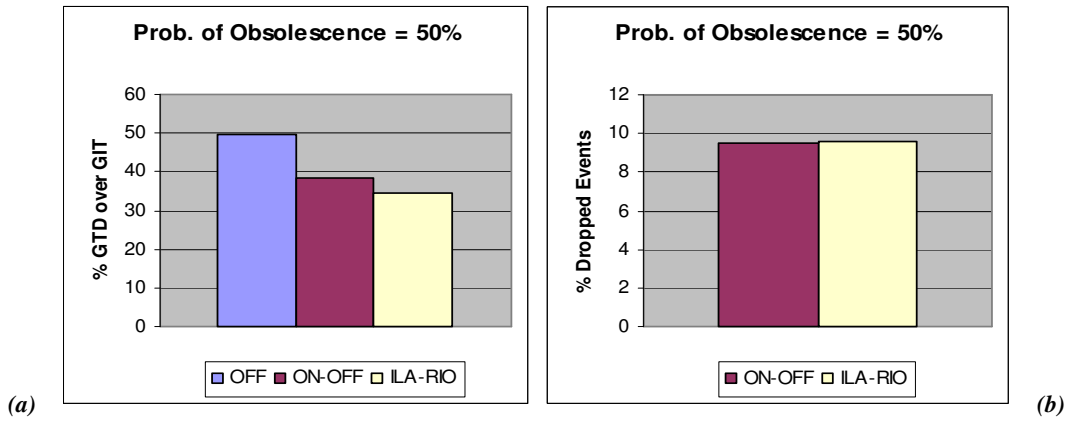


Figure 19. Probability of obsolescence = 50%;
 (a) Event percentage having GTD > GIT; (b) Percentage of discarded events.




Figure 20. Probability of obsolescence = 90%;
 (a) Event percentage having GTD > GIT; (b) Percentage of discarded events.

TABLE V. PERCENTAGE OF OBSOLETE AND VALID DISCARDED EVENTS IN ILA-RIO.

Obsolescence Prob.	50%	90%
<i>Obsolete</i>	9,46%	13,64%
<i>Valid</i>	0,16%	0%

Chapter 4

To Seek the Fairness by Way of the Interactivity

Christopher Columbus's aim, when he sailed for finding Catai (ancient China) lands across the Atlantic Ocean, is perfectly described by the famous claim credited to him: “*Buscar el Levante por el Poniente*”, to seek the East by way of the West. We humbly take inspiration from his genius to synthesize our work in the title of this Chapter. The analogy is represented by the fact that the scheme we propose facilitates fairness by aiming at increasing the interactivity degree in MMOGs.

We demonstrate how it is possible to take advantage of a reduced transmission time to magnify the efficiency of a local lag-type algorithm in ensuring fairness. This represents a very important result since it contradicts the general belief that either interactivity or fairness has to be sacrificed to achieve the other.

1 Exploiting Local Lag Technique

As shortly discussed in Section 2.3 and Section 2.4, Local Lag and other similar algorithms have been proposed to ensure fairness (and consistency) among players in MMOGs [54, 78, 94, 104, 105]. The idea behind this kind of approach amounts to introducing artificial delays in the display of both generated and received game events. These delays are appropriately chosen for each client and depend on their subjective client-server latencies. The aim is that of having each game event simultaneously displayed, after a total amount of time since its creation, on all the players' screens.

Not only does this help in maintaining fairness and consistency, it could also augment the playability of the game. Up to a certain level of injected delay, in fact, players could be more comfortable with a higher but fixed delay than with smaller but variable ones [78]. In the former case, in fact, a player can adapt her/his actions and reaction times to consider also this fixed delay and achieve higher performance than in the latter situation. For instance, car race players will learn how to constantly anticipate the steering when in proximity of a curve; first person shooter players will learn how to aim a little bit ahead of their adversary's position with respect to their moving direction; etc.

Lag compensation can thus be proficiently employed to ameliorate the negative effects of latency on MMOGs. However, the amount of time used as a parameter for local lag generally corresponds to the longest transmission latency experienced by the most unlucky player of the game. In practice, this kind of approach increases game delays and may jeopardize interactivity as in some case the unlucky client may be connected very far away from its server and/or through a slow connection.

Consequently, the efficiency and applicability of the local lag approach strongly depend on the network conditions and on the interactivity degree required by the game. Indeed, especially in the case of a highly interactive MMOG, servers should be optimally located to efficiently serve a large number of customers [76]. Yet, guaranteeing both interactivity and full fairness through local lag can sometimes be achieved only at the cost of impeding the access to some users whose connectivity is irremediably affected by large network delays.

A tradeoff relationship thus exists among scalability (especially in terms of geographical dispersion of the players), interactivity, and fairness. According to this, interactivity and fairness are traditionally seen as incompatible requirements in MMOGs. Conversely, we claim now that upholding interactivity may be useful also to the aim of ensuring fairness. To demonstrate this, we have developed a novel mechanism named Fairness and Interactivity Loss Avoidance (FILA) [164, 165]. Our scheme can be divided into two complementary sub-components. The first one exploits the semantics of the game to drop superseded events and speed up the delivery of game events as seen in Chapter 3. The second one takes advantage of this reduced transmission time to magnify the efficiency of a local lag-type algorithm in ensuring fairness without compromising interactivity.

2 Achieving Fairness through Interactivity

FILA can be thought of as comprised of two complementary parts. The first one, enforced among GSSs, takes substantial inspiration from the aforementioned ILA-RED scheme to speed up the delivery of “fresh” game events by dropping some events which have become obsolete. The second part takes advantage of this reduced transmission time to magnify the efficiency of a local lag-type of algorithm to ensure fairness. FILA utilizes (4) to determine the visualization time of a game event, thus providing fairness without compromising interactivity.

To calculate the appropriate δ in (3), the OD should be determined for each player. For this reason, game events are marked at their creation with a generation timestamp and then sent to the destination: hence, they are orderable. Obviously, a global concept of time has to be maintained in the system. This can be achieved through a variety of solutions that enable the synchronization of GSSs’ physical clocks [67, 111], or by employing new technological synchronization devices such as GPS. Thanks to this, GSSs are able to monitor the OD s of their engaged players and make them available for the FILA algorithm.

More in detail, the first part of FILA drops queued obsolete game events with a certain probability P_d when the average OD , namely $avgOD$, value increases putting at risk the interactivity of the system. The discarding probability P_d is directly proportional to $avgOD$ and dependent on a constant P_{max} . Instead, the value for $avgOD$, at iteration n , is computed through the low-pass filter showed in (17), where w is a parameter that determines how close the average follows the sample trend:

$$avgOD_n = avgOD_{n-1} + w \times (sample_n - avgOD_{n-1}). \quad (17)$$

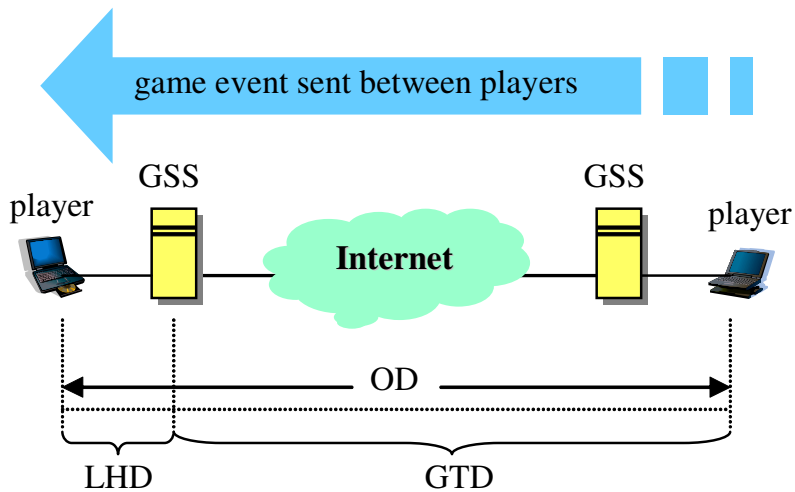


Figure 21. Delay definitions.

More in detail, with FILA, all the game events are regularly processed and forwarded while $avgOD$ is smaller than an alert threshold named $tmin$. When $avgOD$ exceeds $tmin$, the GSSs drop obsolete events with probability P_d , with neither processing nor forwarding them. Finally, if $avgOD$ exceeds the subsequent $tmax$ ($>tmin$) threshold, then P_d is set equal to 1 and all obsolete events waiting for being processed are discarded.

This stabilization mechanism succeeds in reducing $OD_i(e)$ by impacting on $q_i(e)$. In fact, the time spent in queue by a certain event is diminished by the spared processing time of preceding obsolete events which have been dropped with neither processing nor forwarding them. Moreover, since only obsolete events are discarded, FILA fully maintains consistency in the game evolution [62, 109].

To explain FILA more in detail we use the clarifying help of Fig. 21 which provides the graphical definitions for some terms utilized in our explanation: OD , GTD , and LHD (Last Hop Delay).

First of all, it should be noticed that FILA performs its operations on the GSSs. This choice helps us in maintaining a simpler control of the exploited game platform. Under that circumstance, however, for each event e , GSSs can compute $GTD(e)$ but not $LHD(e)$. An estimation of $LHD(e)$ is necessary in order to compute OD and utilize it in our algorithm. For this reason, each GSS continuously monitors the latencies to each of its engaged players and maintains a variable named λ_{GSS} . The value of this variable represents the maximum among the latencies from the considered GSS to each of its connected clients (this set of clients is named C_{GSS}) and is calculated as follows:

$$\lambda_{GSS} = \max_{i \in C_{GSS}} \{ LHD_i \}. \quad (18)$$

However, we cannot let some irremediably delay-affected client to excessively impact on the calculations performed by our scheme. Utilizing, in FILA, the excessively high λ_{GSS} generated by some player connected very far away from the GSS, in fact, would result in very high sample (and $avgOD$) values with respect to GIT . In this case, FILA would increase the aggressiveness of its discarding function as perceived by all the players with no positive results: the “unlucky” player will still not be able to receive game events with delays below the interactivity threshold. For this reason, we need to consider a *Delay Upper Bound (DUB)* in order to limit the impact of “unlucky” players on its algorithm. To this aim, (19) provides the formula for a fundamental parameter utilized by FILA to handle the impact of $LHD(e)$ on its algorithm:

$$\sigma = \min \{ \lambda_{GSS}, DUB \}. \quad (19)$$

The usage of this parameter depends on the employed version of our scheme and is explained in Section 5. To determine DUB we rely on a heuristic that dynamically computes its value based on the general network condition during the game. Its formula is as follows:

$$DUB = GIT - \max \{ GTD \}. \quad (20)$$

where $\max\{GTD\}$ represents the largest among the $GTDs$ experienced over all the connections between each GSS and the players engaged by the other GSSs.

To compute DUB , each GSS has hence to periodically determine the GTD that features in average the slowest of its connections with players engaged by the other servers. Then, this value has to be communicated back to all the other peers in order to allow a global knowledge of the worst GTD endured by each GSS. Finally, the highest among these maximum $GTDs$ can be univocally determined by each of the GSSs and used to determine the global DUB in the system as shown in (20).

The second part of FILA is simply in charge of equalizing the delay differences among players with a local lag-type scheme that appropriately computes the δ value shown in (3) so as to satisfy (4) whenever possible. We are going now to empirically demonstrate how the combination of phase one and two is effective in ensuring fairness and interactivity while allowing a scalable number of contemporary players.

We compare the regular local lag (LL) mechanism against FILA. In particular, LL embodies the traditional local lag scheme with no discarding mechanism for obsolete events. Even in this case, however, as for all the other compared protocols, the algorithm is not allowed to introduce artificial delays if this would result in jeopardizing interactivity (i.e. $t^{v(e)}$ cannot be set greater than $t^{g(e)} + GIT$).

Focusing on FILA, we have set $t_{max} = GIT$ and $t_{min} < t_{max}$. Moreover, the estimation of the $LHD_i(e)$ is performed through (18) and (19) and utilized to calculate the new sample upon every new event arrival as follows:

$$sample(e) = GTD(e) + \sigma. \quad (21)$$

3 Simulation Assessment

It is well known that MMOG service providers should appropriately position their game servers in such a way that their target player market would be located within a circle having 150-180ms of latency diameter [76]. Following this rule and aimed at creating a configuration able to factually support a highly interactive MMOG, we have simulated a constellation of five GSSs deployed across U.S.A. by choosing optimal market locations.

Clients are supposed to be distributed all over the North American continent and connected through various access technologies that provide them with different access delays. We have focused our attention on the event receiving aspect of a single GSS (GSS_0), pretending that the other GSSs are sending events to it (without any loss of generality).

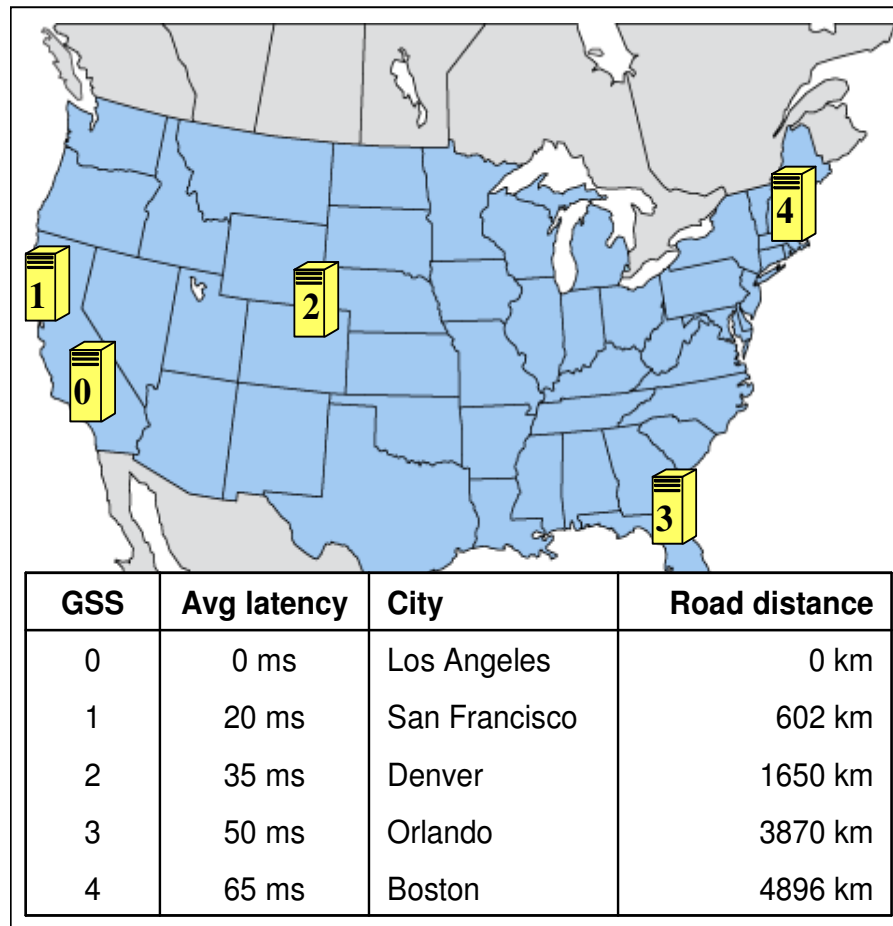


Figure 22. Game servers deployment.

Inspired by literature [77], the GTD values were chosen based on a lognormal distribution whose approximate average was obtained by means of repeated runs of the ping application. More in detail, game events coming from clients connected to the sending GSSs (i.e. GSS_1 – GSS_4) and traveling towards GSS_0 experience average latencies as reported in Fig. 22, with a standard deviation of 10ms. Further, several scenarios were considered where the values of $\max_{i \in C_GSS} \{ LHD_i \}$ were chosen for each GSS

within the following set {25ms, 50ms, 75ms, 100ms, 125ms, 150ms}. This choice is justified by the consideration that clients should be located within a circle having a maximum latency diameter of 150ms. We assumed to have 10 clients connected to each GSS, engaged in a fast-paced game, and generating a new action every 300ms in average. This results in a flow of game events having 30ms of inter-departing time. Finally, the average game event size (200 Bytes) was inspired by literature about games as well [17].

Focusing on the parameters in the FILA algorithm, we have set $w = 1/8$ for all the simulations. The alert threshold $tmin$ was equal to $GIT - 100ms$. Moreover, the probability that an event makes obsolete preceding ones was set to 90%. This represents a realistic scenario for a vast plethora of possible games (e.g. adventure, strategic, vehicle race, flight simulator, etc.), where most of the events are just independent movements as seen in Section 3.6.

Each experiment was identically replicated to compare the outcomes of FILA against the regular LL algorithm. In [104], Zander *et al.* demonstrated that there is a statistically significant difference between the mean kill rates of player groups which are affected from diverse client-server latencies. In essence, lower latencies results in higher mean kill rates and thus in unfairness. Coherently, we have chosen to evaluate as a performance parameter the percentage of events that were delivered by GSS_0 to *all* of its players in time to be contemporary visualized before the GIT expiration. We are hence considering the achievement of per-event interactivity and fairness.

4 Results

In this Section, we demonstrate through results how FILA is able to ensure a higher interactivity and fairness degree if compared to the traditional LL scheme. Moreover, we provide a scalability evaluation of the two mechanisms showing how FILA improves its performance in situation with intense traffic and outperforms regular LL.

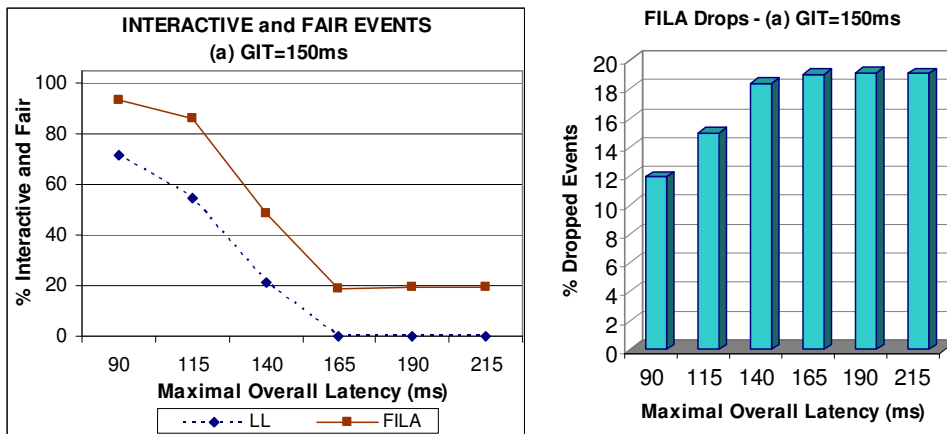


Figure 23. Interactivity and fairness improvement (left) and dropped events (right) with $GIT=150ms$ and $AIDT=30ms$.

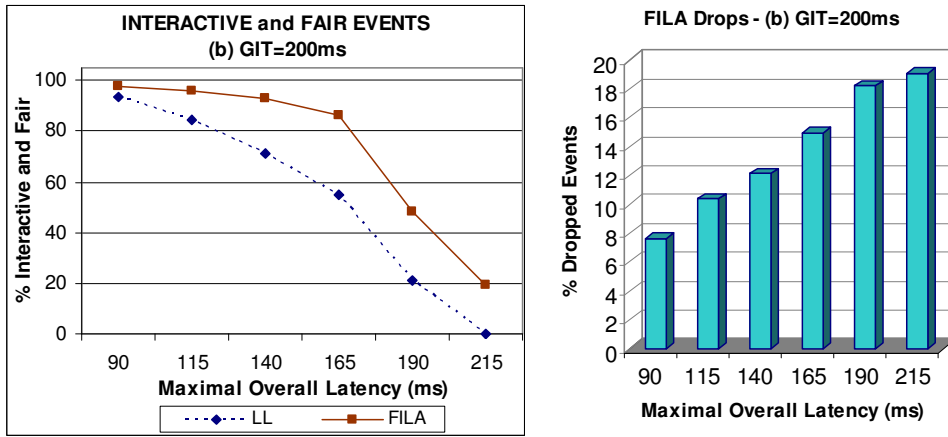


Figure 24. Interactivity and fairness improvement (left) and dropped events (right) with GIT=200ms and AIDT=30ms.




Figure 25. Interactivity and fairness improvement (left) and dropped events (right) with GIT=250ms and AIDT=30ms.




Figure 26. Interactivity and fairness improvement (left) and dropped events (right) with GIT=300ms and AIDT=30ms.

4.1 Interactivity and Fairness

Fig. 23, Fig. 24, Fig. 25, and Fig. 26 show, respectively, four different sets of experiments, obtained varying the *GIT* from 150ms to 300ms. Each set was comprised of six different experiments and each experiment consisted in the transmission of about 4000 game events which experienced, in the worst case, a maximal overall latency whose value is reported on the x-axis of each provided chart. The maximal overall latency represents the largest average latency experienced on the connection between any two players in the system.

The leftmost graphs of Fig. 23, Fig. 24, Fig. 25 and Fig. 26 show the percentage of game events that GSS_0 was able to deliver to all of its engaged players in time to be simultaneously delivered with an *OL* lower than *GIT*. It hence represents the amount of events which satisfied condition (2) and were thus fairly processed by all the clients.




Figure 27. Interactivity and fairness improvement (left) and dropped events (right) with $GIT=150ms$ and $AIDT=20ms$.

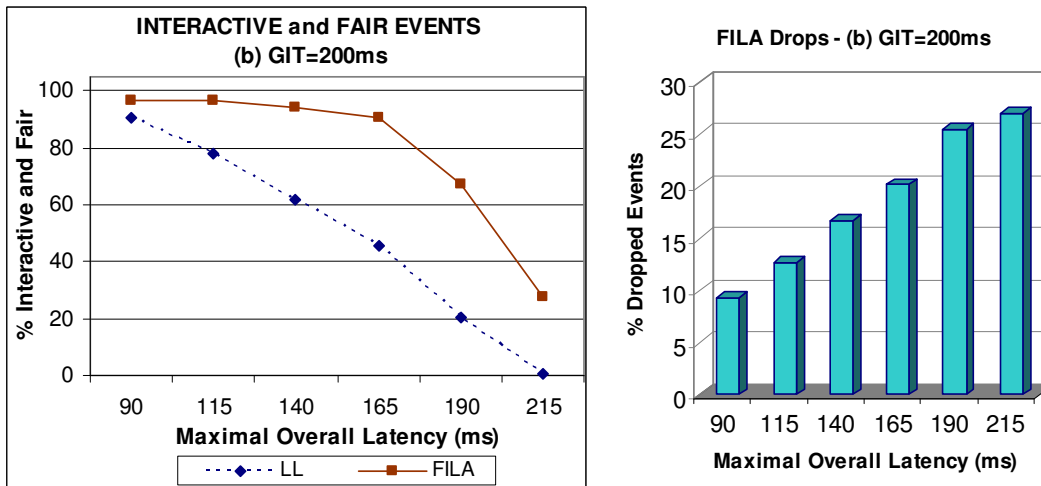


Figure 28. Interactivity and fairness improvement (left) and dropped events (right) with $GIT=200ms$ and $AIDT=20ms$.

As can be seen from these graphs, having a higher *GIT* improves the efficacy of both the evaluated schemes since larger local lags can be utilized. However, regular LL algorithm experiences a premature performance decrease when the maximal overall latency increases even if it is still far from the *GIT*. Instead, FILA ensures a good fairness degree for a larger set of overall latencies.

Obviously, in those configurations where the maximal overall latency is close to (or surpasses) *GIT*, both schemes cannot overwhelm network conditions, thus achieving poor fairness (and interactivity). Even in this case, however, FILA behaves better than the regular LL algorithm.

FILA pays these better results with the drops of some obsolete events. Specifically, the rightmost charts of Fig. 23, Fig. 24, Fig. 25, and Fig. 26 reveal the percentage of game events which were discarded by FILA. In all the considered cases, less than 20% of the game events were dropped and these events were exclusively obsolete ones.

Results turn out to be even better if we focus only on those cases where the overall latency is not irredeemably high with respect to *GIT*. Considering the configurations when the maximal overall latency is lower than *GIT* by 35ms or more, we find that FILA always guarantees more than 86% of fairly delivered game events with less than 15% of dropped events.




Figure 29. Interactivity and fairness improvement (left) and dropped events (right) with *GIT*=250ms and *AIDT*=20ms.




Figure 30. Interactivity and fairness improvement (left) and dropped events (right) with *GIT*=300ms and *AIDT*=20ms.

4.2 About Scalability

In order to test the scalability of FILA and LL, we have decreased the AIDT to generate scenarios with a higher level of game traffic in the network. In particular, Fig. 27, Fig. 28, Fig. 29, and Fig. 30 refer to configurations with 20ms of AIDT, while Fig. 31, Fig. 32, Fig. 33 and Fig. 34 correspond to the cases where AIDT is equal to 10ms. Again, in each figure we present the outcomes for four different GIT values: 150ms, 200ms, 250ms, and 300ms, respectively. The leftmost charts show the percentage of game events that were fairly and interactively delivered to all the clients engaged by GSS_0 . Instead, the rightmost ones reveal the percentage of game events which were discarded by FILA.

As one can expect, the higher the game traffic, the lower the interactivity and fairness degree provided by LL. On the contrary, not only is FILA able to manage higher traffic, but its performance actually improves when the AIDT decreases. This surprising result has a simple explanation. Higher rates in game event transmissions result in larger queues at GSS_s ; these queues contain packets that have not yet been processed. This represent an insurmountable problem for LL since $q_i(e)$ increases for all clients putting at risk the performance of the system without having any countermeasures. With FILA, instead, a larger queue of game events at a certain GSS represents also a resource. In fact, obsolete game events in queue can be discarded, thus reducing the $q_i(e)$ that a subsequent event e will experience in its traveling towards the various clients i .

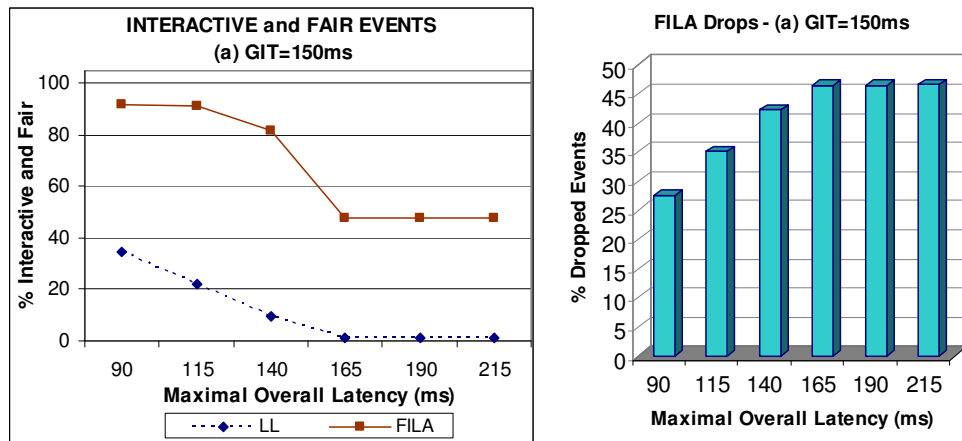


Figure 31. Interactivity and fairness improvement (left) and dropped events (right) with GIT=150ms and AIDT=10ms.

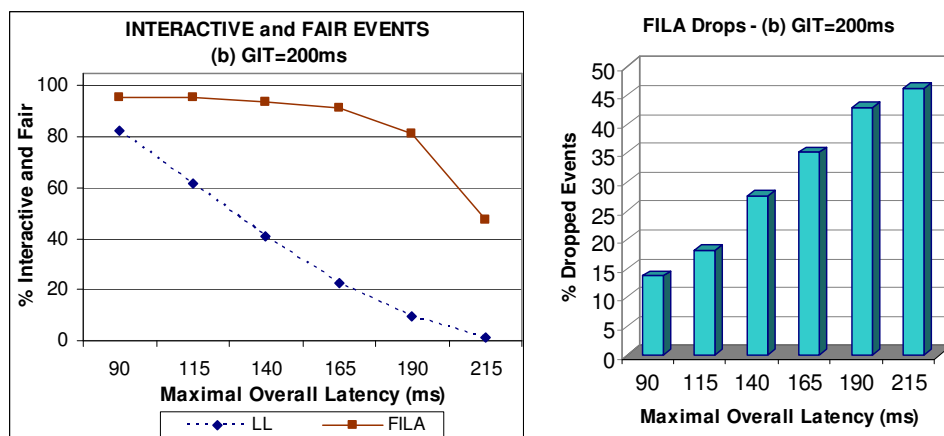


Figure 32. Interactivity and fairness improvement (left) and dropped events (right) with GIT=200ms and AIDT=10ms.

As a proof for our rationale, we can notice that the number of obsolete game events dropped by FILA increases when decreasing the AIDT. This is caused by higher *avgOD* values due to the increased traffic, but is also possible thanks to the presence of more game events in queue that FILA can exploit to drop obsolete ones.

Finally, analogously to the scenario with 30ms of AIDT, even when the AIDT is set equal to 20ms or 10ms, the percentage of discarded game events remains still reasonably small.




Figure 33. Interactivity and fairness improvement (left) and dropped events (right) with GIT=250ms and AIDT=10ms.




Figure 34. Interactivity and fairness improvement (left) and dropped events (right) with GIT=300ms and AIDT=10ms.

Chapter 5

Wireless Home Scenario

The market is currently heading toward houses where all the devices (e.g., computers, televisions, phones, intelligent appliances, etc.) will be wirelessly connected to the home network and possibly controlled by a single hub. This convergence point might be represented by the media center which will expand its features becoming, within few years, the engine of the home network and control the whole home connectivity.

In this context, take, for example, a mid-class American household where a family of four people lives: two teenage kids and the hardworking parents. Each family member presumably owns several networked personal portable devices such as PDAs, MP3 players, game consoles and digital cameras; all these being also connected to the home network.

Based on the market trends, we consider that all those devices are wirelessly connected to a media center that controls the in-house media distribution and provides access to the Internet as well as to the cable television and companies providing external services (e.g., the alarm company).

Moreover, we also assume that several family members will be accessing the household network at the same time according to their work or leisure needs. In particular, for the sake of our study, we consider the following family scenario:

- i) one teenager is watching a movie, streaming it from the media entertainment center;
- ii) the other one is playing with his latest MMOG against a crowd of buddies across the Internet;
- iii) the father is having a conversation through an IP based video-chat;
- iv) the mother is downloading the last U2 greatest hit compilation from the Apple iTunes© music store.

In the above everyday life picture it is worth noticing that each of the aforementioned employed applications features different requirements in terms of network performance, as well as suffer from very specific problems all due to the best effort nature of the Internet transport protocols.

These are as follows:

- **Video-Streaming.** Streaming applications are affected by the jitter phenomenon while are resilient to some packet loss; a network designed mainly for video-streaming should minimize the jitter. Buffering techniques can be applied to minimize the impact of variable network conditions.
- **Video-Chat and Massive Multiplayer Online Games.** Both this applications require a high degree of interactivity, they greatly suffer from delays and packet jitter while may tolerate some packet loss, as we have seen for the case of superseded events.
- **iTunes Music download.** A music download activity is typically performed using TCP, hence this type of application is resilient to jitter and delays but decreases the sending rate in presence of losses: it hence does not tolerate any error losses (losses that do not depend from congestion).

1 Transmission Control Protocol

The success of the Internet is based on several factors; among these, one of the most important is the ability to provide a reliable medium for information exchange and file downloading. In particular, traffic control functionalities for the most common applications (i.e., FTP, HTTP, SMTP, Telnet) are provided by the Transmission Control Protocol (TCP). TCP was initially designed to provide an end-to-end, connection-oriented, and reliable service in the ARPANET [138], and later, in the Internet. TCP addresses two major issues: reliability and congestion control [139]. To achieve the second goal, TCP adapts the sending rate to avoid network overflow or falling into service starvation. TCP congestion control has been studied by the research community for the last 25 years, leading to several TCP variants with and without the explicit intervention of the network layer (a survey can be found in [140]).

The most popular version, TCP New Reno, implements a congestion control algorithm, known as the AIMD (Additive Increase, Multiplicative Decrease) algorithm. In this context, the *sending window* utilized by the sender represents the number of packets that the sender can send towards the destination without having yet received any acknowledgment about their delivery from the receiver. Indeed, for every packet received, the receiver sends back to the sender an acknowledgment (ack) that identifies the next packet expected and implying that all the precedent packets in the sequence were successfully delivered. Acks (and *Time Outs* [139]) are also used to determine packet losses and to communicate the *advertised window* back to the sender. The advertised window provides the sender with a limit to the maximum transmissible sending rate. The final sending window, in fact, is the minimum between the advertised window and a variable named *congestion window*.

The very basic concept of the congestion control scheme can be summarized as follows:

- The number of packets sent out without having yet received back their corresponding acks cannot be higher than the current sending window.
- For each successfully delivered packet, a new one is sent.
- When a whole congestion window of packets has been successfully delivered, the congestion window value is increased by 1.
- When a packet loss is detected the TCP sender assumes congestion on the path to the receiver and decreases its congestion window by half (or to 1 if a Time Out occurred).

This scheme has been developed following the end-to-end paradigm by which the two involved end nodes do not have any explicit information about the links connecting them. In essence, the Internet is seen as a *black box* whose contents remain unknown and all the intelligence is left at the edge. Sender and receiver are unaware of the available bandwidth on the links among them and of the possible presence of other flows along the same path. The sender has hence to continuously probe the channel to make use of bandwidth that might be available and to back off when congestion is detected.

2 The IEEE 802.11

As the availability of digital entertainment devices increases rapidly, the need for interconnecting them is felt as ever more urgent, as well as the necessity to extend the reach of entertainment centers to the wireless domain. In the today's market, IEEE 802.11 based wireless LANs are *de facto* emerging as the candidate to lead the mobile revolution providing wireless connectivity and advanced functionalities in terms of flexibility, security and throughput to support entertainment applications ranging from networked games to in-house digital audio/video distribution and live conferencing etc. [141]. In this context, technical standards are currently being defined that address both the methods of wired/wireless interconnection and the means to guarantee a full interoperability between digital entertainment appliances. Yet not much work has been done in the direction of understanding how can the Internet native language (i.e. the TCP/IP protocol) take over this complex scenario for efficiently delivering

digital contents to entertainment devices, and which is the impact of diverse MAC layer settings over the Internet native transport protocols (e.g. TCP, UDP) during the distribution of in-house entertainment.

Recalling its main characteristics, the 802.11 MAC layer protocol attempts to face the packet loss problem by implementing its own retransmission scheme [142]. In particular, lost packets are retransmitted after a certain period of time without having received any corresponding ack. Successive retransmissions for the same packet are repeated up to a maximum number of time, which is by default set to 4 in the standard IEEE 802.11, or until receiving a successful ack. A back off mechanism determines the retransmission timeouts. This scheme hides wireless error losses from the TCP's congestion control mechanism, thus avoiding deleterious multiple reductions of the data sending window. On the other hand, local retransmissions affect packet delivery delay by increasing its variability and thereby affecting time-constrained applications such as audio or video-stream.

We are aimed at studying several realistic wireless networking scenarios for entertainment over IEEE 802.11 networks. In particular, we are going to consider the IEEE 802.11g standard since its wide diffusion and its increased link capacity (54Mbps) with respect to IEEE 802.11b (11Mbps). In essence, among the MAC protocols currently in commerce, IEEE 802.11g represents the best candidate for being factually exploited in wireless homes to support multimedia applications.

3 Queuing Delay

As shown in (5), several delay components determine the final delivery time of each game events and a significant one is represented by queuing time. Queues are built up along the path from sender to receiver when the arriving rate of events at a certain node is superior to the serving rate featuring that node. For instance, there could be routers along the path that receives more packets per unit time than the transmitting rate available on the outgoing link on which those packets have to be forwarded. Another example is represented by Game servers that receive more game events per unit time than the rate at which they are capable of processing them.

However, as recently demonstrated by measurements on a real OC48 link, the capacity of the Internet is generally larger than the aggregate bandwidth utilized by transiting flows [156]. Moreover, more and more providers are offering today *guaranteed high speed connectivity* to home customers [150, 151, 152, 153, 154, 155]. In essence, tools are offered to customers to verify that their connection is actually able to support as much traffic as the declared bandwidth. This implies that the bottleneck of the connection is generally located at the edge of the path connecting sender and receiver.

Focusing on MMOG deployment, we can further support this assumption. In fact, revenues for MMOG providers come from the subscription payments of many satisfied customers. Therefore, every commercial MMOG is generally supported by adequate resources in terms of connectivity speed among, number and capability of their servers [97]. Moreover, we have just demonstrated in Section 3 and Section 4 how efficient synchronization policies among GSS can be put to good use to improve interactivity and fairness degree.

However, even when the game network platform is able to bring game events to our house with delivery times within the *GIT*, still problems may arise in the last hop, which represents the bottleneck in terms of the available capacity for the connection. In fact, it might be the case when the Access Point (AP) receives packets at higher rates that it can forward them to destination. This can happen for several reasons as, for instance, the fact that the wireless medium allows the transmission of only one packet at a time and is not full-duplex as wired links.

Moreover, interference, errors, fading, and mobility may cause packet losses which are handled by the MAC protocol through local retransmissions. These local retransmissions hide error losses to the TCP and are useful to increment the reliability of the connection. Without them, the TCP would misinterpret error losses as congestion evidences and reduce its sending rate decreasing its performance. On the other hand, retransmissions follows the well known back off mechanism by which an increasing amount of time is utilized to determine whether a packet has been lost and hence retransmit it. The 802.11 MAC protocol performs up to seven retransmissions of short packets (i.e., RTS/CTS, acks) and four retransmissions of long packets (i.e., data packets) [157]. This means that following packets have to wait in queue until the preceding one or one of its retransmissions finally reaches the receiver and the corresponding ack is successfully sent back.

Finally, the same wireless connection might be shared by several devices and applications that increase the congestion level causing queuing. As it is well known, TCP connections have an aggressive behavior and continuously probe the channel for more bandwidth until queues are fully utilized and overflowed. In presence of persistent TCP connections it is hence very likely to happen that queues are steadily fully utilized, thus periodically slowing down the delivery time of each packet, and deteriorating the performance of time-sensitive applications such as MMOGs.

At the same time large buffers helps TCP-based flows in keeping a high sending rate. This happens for several reasons but the most important ones are: i) the link successive to the buffer remains fully utilized for longer periods of time since there are (almost) always packets in queue, ready to be sent as soon as possible, and ii) traffic bursts can be more easily accommodated thus reducing packet losses and maintaining higher sending rates for longer periods of time. In essence, a tradeoff relationship exists among the per-packet delay and the total goodput achieved. The solution for this tradeoff depends on the queue size and on its utilization.

As we demonstrate in Section 5.6.1, delay increments caused by TCP-based traffic could hit also tens of milliseconds, which represent a huge waste of time when trying to deliver game events within a *GIT* of 150ms.

4 Proposed Solutions

To solve the aforementioned problem we propose different possible solutions. In particular, we may divide them into two main classes depending on the correspondent networking layer where they are implemented: transport layer solutions and MAC layer solutions.

The former mainly regards exploiting some of the existing features of regular TCP or employing an alternative TCP version. The solutions belonging to this class does not necessary rely on the fact that the last link will be a wireless one. The latter, instead, involves modifications of the 802.11 MAC layer and thereby it is specifically intended for the wireless media [166].

For all the proposed solutions we investigate both their efficiency and factual deployability to expose pros and cons and conclude with a winner.

4.1 IEEE 802.11 Parameters Setting

The first proposal regards the utilization of more appropriate setting for parameters of the IEEE 802.11 MAC protocol. Parameters such as maximum number of retransmissions and buffer size were, in fact, decided in a period when the TCP-based traffic was largely predominant in the Internet. The main concerns for designers were hence reliability and high throughputs.

Nowadays, UDP-based real time applications are becoming more and more popular hence demanding for low delays in packet delivery. This kind of applications are resilient to some packet loss whilst cannot tolerate delays in packets delivery. For this reason, it is preferable to drop a packet than to wait time in retransmissions.

This obviously contradicts the initial assumption that reliability is the most important issue over wireless links. Therefore, 802.11 parameters should be modified to make it more sensitive towards real-time application needs. In particular, the number of local retransmissions could be diminished in order to find an efficient solution between reliability and low delays in packet delivery.

In the same way, large buffer sizes at the AP help TCP connection to maintain large sending rate for longer periods and diminish the impact of burst traffic. On the other hand, to a larger buffer corresponds a longer queuing time experienced when the buffer is full, thus jeopardizing the performance achieved by time-sensitive applications. By adjusting the buffer size to an appropriate value we can again try to find an optimal compromise between the needs of TCP-based applications and real-time applications.

4.2 TCP Vegas

TCP Vegas embodies one of the most cited alternatives to regular TCP New Reno in scientific papers. Its main point of interest is represented by the fact that TCP Vegas tries to avoid congestion instead of blindly increase its sending window until a packet loss occurs; it hence perfectly fits our needs for low buffer utilization.


```

for every RTT
{
    if (cwnd/RTTmin - cwnd/RTT) < α then cwnd++
    if (cwnd/RTTmin - cwnd/RTT) > β then cwnd--
}

for every loss
    cwnd := cwnd/2

```

Figure 35. Pseudocode of TCP Vegas congestion control.

Indeed, while TCP New Reno utilizes packet loss to determine network congestion, TCP Vegas is sensitive to end-to-end queuing delay. With TCP Vegas the sender monitors the difference between its expected rate and the actually achieved one. The difference is compared to a couple of parameters, namely α and β , to determine whether the congestion window has to be incremented or decremented (by 1) in the next RTT [143, 144]. The parameter α and β determine the amount of buffer each flow is permitted to occupy.

More in detail, each TCP Vegas source estimates the number of its own packets in the buffer by monitoring the RTT. It then accordingly adjusts its congestion window to maintain its estimated rate between the two predetermined parameters α and β . If the difference between the expected rate and the achieved one is smaller than α , then the congestion window is increased by 1. If this difference surpasses β , the congestion window is decremented by 1. Finally, if the difference is comprised between α and β then the sender rate is in its stability region and no particular operation is performed.

The pseudocode of the congestion avoidance algorithm of TCP Vegas is reported in Fig. 35 where cwnd is the congestion window. The pseudocode also shows that, in case of a packet loss, the congestion window is halved.

If the buffer at the bottleneck is large enough then TCP Vegas reaches equilibrium. In this case, TCP Vegas flows should experience zero packet losses, a stable congestion window, and a queue which is proportional to the number of TCP Vegas flows present. Otherwise, it oscillates like TCP New Reno flows. Finally, TCP Vegas has been proven to fairly share the link with other TCP Vegas flows but it behaves too conservatively in presence of regular TCP flows (i.e., New Reno, Sack). Regular TCP, in fact, fully exploits the available buffer space and TCP Vegas interprets the consequent RTT trend as an indicator of excessive congestion, thus progressively reducing its sending rate to very low values [148].

TCP Vegas represents a very appealing transport protocol with interesting features such as the possibility to have some control on the amount of buffer utilized. This, in fact, results a fundamental property when having to efficiently deal with contemporary real time traffic. On the other hand, the dramatic efficiency decrease experienced when competing with regular TCP traffic makes unfeasible its factual deployment.

4.3 Limited Advertised Window

We are aiming at finding the best solution to the tradeoff relationship existing between TCP throughput and real time application delays. Moreover, the two types of traffic should be able to coexist without interfering each other and the employed solution should be easily and factually deployable.

Starting from the last point, i.e. the deployability, it is evident how a technique that would exploit existing features of the already utilized protocols could be easily implemented in a real scenario. Moreover, we deem that an optimal tradeoff between throughput and low delays could be achieved by maintaining the sending rate of the TCP flows high enough to efficiently utilize the available bandwidth but, at the same time, limited in its growth so as to not utilize buffers.

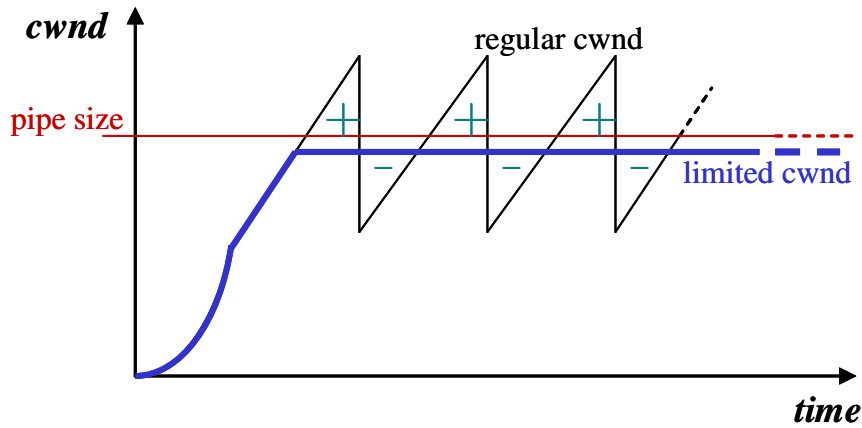


Figure 36. Comparison between regular and limited sending windows.

In this way, in fact, the throughput is maximized by the absence of packet losses which would halve the congestion window, while the delay is minimized by the absence of queues.

To better understand how limiting the congestion window could guarantee the same or even a higher throughput with respect to utilizing regular TCP, we show in Fig. 36 a typical saw tooth shaped sending window of a regular TCP and overlap it with a limited one. As it is evident, the latter is more stable since it does not use the buffer at the bottleneck link and consequently experiences no losses. The minus signs in the chart represent situations in which the regular congestion windows provides TCP with a sending rate which is inferior to that guaranteed by the limited congestion window, while the plus signs represent the inverse situation (generally accompanied by having packets queuing on the buffer corresponding to the bottleneck link). If the upper bound for the congestion window is appropriately chosen, the balance between the plus and minus signs will guarantee to the limited congestion window an equal or even superior final throughput with respect to the regular congestion window, whilst avoiding queuing delays.

To achieve this desirable result we need first to address two important issues: how to determine an appropriate upper bound and how to apply it in practice to the sending window.

Regarding the first point, the most appropriate formula can be derived from the two main goals we want to achieve: i) full utilization of the available bandwidth and ii) no queue delays. Real time traffic generally exploits UDP and this transport protocol has no congestion control mechanism. Some smart applications, however, implements some sort of congestion control at the application layer [158, 159]. In any case, to avoid queue delays, the aggregate bandwidth utilized by TCP flows cannot exceed the total capacity of the bottleneck link diminished by the portion of the channel occupied by the concurrent real time traffic.

In essence, the maximum sending rate for each TCP flows at time t , namely $TCPubrate(t)$, is represented by:

$$TCPubrate(t) = \frac{(C - UDPtraffic(t))}{\#TCPflows(t)} \quad (22)$$

where $UDPtraffic(t)$ represents the amount of bandwidth occupied by UDP-based traffic at time t , $\#TCPflows(t)$ is the concurrent number of TCP flows, and C corresponds to the capacity of the bottleneck link.

The second issue we need to address is how to practically employ this formula in order to have it working in a real scenario. This means i) determining an effectively deployable way to utilize it with the current state of the art of the Internet, ii) identifying the location for its implementation, and iii) proposing a method to compute the value of the various variables.

To solve the first issue, we have to limit the required scope of intervention since modifying the whole Internet in order to run our scheme would obviously not be a feasible option. Moreover, it would

definitely be better if we could make good use of some feature already present in the regular TCP. For this reason we propose to exploit the existing advertised window to limit the TCP sending rate. In fact, as seen in Section 5.1, the actual sending window is determined as the minimum between the congestion window and the advertised window. The advertised window perfectly embodies an upper bound to the congestion window and is already implemented in all TCP versions. By appropriately modifying it, we can achieve both efficiency and low delays.

The advertised window is generally determined at the receiver; however, this could not represent the most suitable place for the modification we need to perform. Indeed, to determine the most appropriate value for the advertised window, we need a comprehensive knowledge about all the flows that are transiting through the bottleneck (i.e., the last hop links). Since all the flows have to pass through the AP, this represents the node where we could be able to implement our scheme. The AP may also coincide with the Media Center in a wireless home and the mechanism can take advantage of this. In particular, by spoofing the transiting traffic at the AP and/or utilizing the information hold by the Media Center, we can obtain all the information required. In any commercial Operating System it is possible to know which kind of connection is in use and which its nominal speed is just by looking at the status of the network interface. Through spoofing the channel or exploiting information known at the Media Center we can also infer the number of active TCP connections and the aggregate amount of current UDP traffic. The AP can hence easily compute the best $TCP_{ubrate}(t)$ utilizing (22) and modify the advertised window included on the transiting acks accordingly.

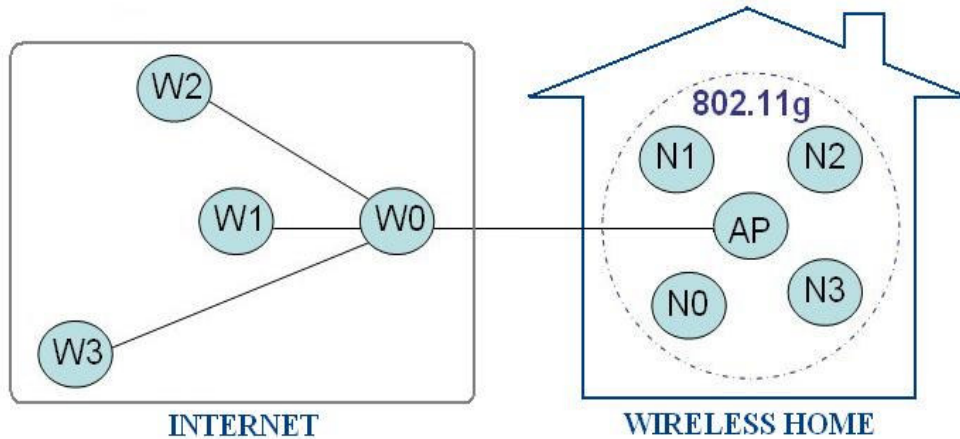


Figure 37. Simulated topology.

TABLE VI. SIMULATION CONFIGURATION OF THE WIRED LINKS.

Node 1	Node 2	Physical Latency	Link Capacity	Queue Size
W1	W0	10ms	100Mbps	140pkts
W2	W0	20ms	100Mbps	140pkts
W3	W0	30ms	100Mbps	140pkts
W0	AP	10ms	100Mbps	140pkts

5 Simulation Assessment

In order to analyze in depth our scenario, we have utilized the well known NS-2 network simulator (version ns-2.28) [1]. Our adopted configuration of the nodes and links can be easily visualized by the means of Fig. 37. In particular, the house environment is represented by four mobile nodes named N1, N2, N3 and N4, and the Media-Center that incorporates also the AP. The MAC layer parameters have been set accordingly to the IEEE802.11g standard. The simulation outcomes showed us that we were able to reach a maximum achievable bandwidth of circa 20Mbps. This represents a reasonable value over the declared 54Mbps even in the real world [174].

Focusing on the wired links, their one-way delays and capacities have been configured as listed in Table VI, while their queue sizes have been set equal to 140 packets. This value comes out by multiplying the longest RTT with the smallest link capacity on the path (i.e. the bottleneck) which is represented by the 20Mbps effectively available over the wireless link.




Figure 38. In-home wireless scenario.

TABLE VII. SIMULATED APPLICATION FLOWS.

From	To	Flow Type	Transp. Prot.	Start	End
AP	N0	Video-Stream	UDP	0s	180s
W1	N1	Online Game	UDP	45s	180s
N1	W1	Online Game	UDP	46s	180s
W2	N2	Video-Chat	UDP	90s	180s
N2	W2	Video-Chat	UDP	91s	180s
W3	N3	FTP	TCP	135s	180s

As shown in Table VII and Fig. 38, several kinds of applications have been run over this network topology. In order to uplift the trustworthiness degree of the simulations, we have exploited real trace files for the video-stream and for the video-chat. Specifically, adopted trace files correspond respectively to high quality MPEG4 *Star Wars IV* for the movie, and two VBR H.263 Lecture Room-Cam for the video-chat, as can be found in [145].

The parameters characterizing the game-generated traffic have been chosen following the directions provided by scientific literature in this field. Indeed, we can assume that the user in the house is engaged in one of the very popular first person shooter games, e.g. Quake Counter Strike, with other ~25 players, geographically away from each other and connected through the Internet. Hence, to model the traffic generated by this kind of MMOG (packet size and interarrival time), we can use some of the approximations suggested in [146], which are based on real game platform measurements.

In particular, in our simulation game events have been generated at client side every 60ms; while the server was transmitting game state updates every 50ms toward the client. Moreover, packet size has been set to 42Bytes and 200Bytes, respectively for client and server generated game packets.

Simulation experiments have been replicated to examine the effects generated by differently setting some of the parameters involved in the scenario. Table VIII lists all the variable parameters in the simulations; each combination of their possible values has been simulated. In particular we have tried several values for long packet retransmissions going from the regular 4 down to 1. The distance between the AP and the mobile device was either 5 or 10m which represent two realistic values for a normal house. The buffer size at the AP was set either equal to 50 or to 100 packets, as these are the two most common values in commerce.

In particular, it is worth to conclude this discussion by mentioning our experimental choices with respect to the Shadowing Model, which is a realistic and widely utilized signal fading model available in NS-2. We followed the directions provided by the official NS-2 manual to represent a home environment partitioned into several rooms. Specifically, in our simulations, the path loss exponent of the Shadowing Model was always set equal to 4, while the shadowing deviation had alternatively the value of 7 and 9. Transmitted signal attenuation grows with the increase of these parameters; we hence expect to face higher percentage of packet losses over the wireless media when setting the shadowing deviation to 9.

However, where not differently stated, simulations were run utilizing some realistic default values for the simulative parameters listed in Table VIII. In particular, we had:

- buffer size at the AP = 100 packets;
- distance between the AP and the mobile device = 10m;
- max number of MAC retransmissions for long packets = 4;
- shadowing deviation of the shadowing model = 7.

TABLE VIII. CHANGING PARAMETERS IN THE SIMULATED CONFIGURATIONS.

Parameter	Values	Comment
MAC data retransmissions	1, 2, 3, 4	default value = 4
shadowing deviation	7, 9	medium, high
user-AP distance (m)	5, 10	same room, different room
MAC queue size (pkts)	50, 100	common default values




Figure 39. Example of online gaming interarrival delays.




Figure 40. Example of online gaming jitter.

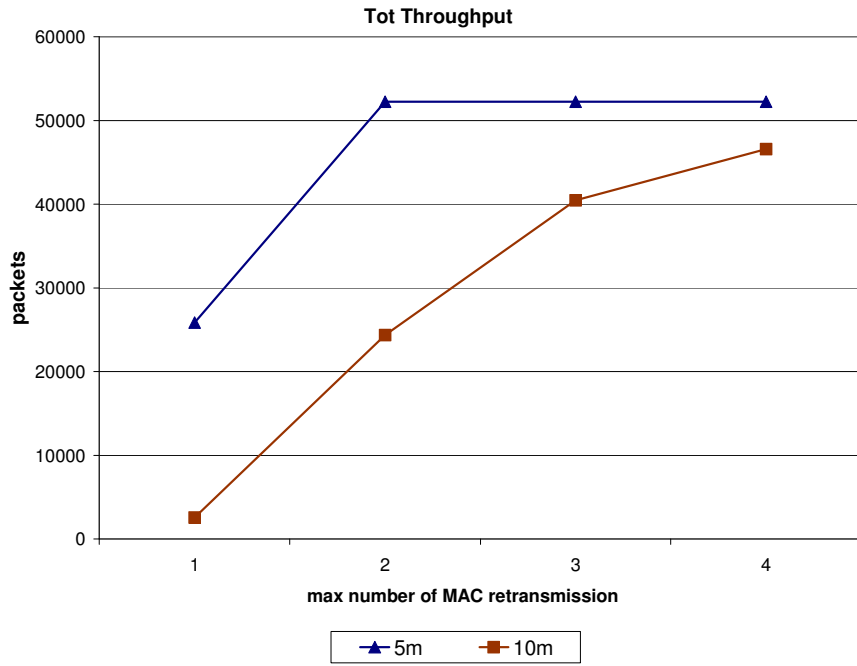


Figure 41. FTP total throughput with different user-AP distances; shadowing deviation = 9, MAC queue size = 50pkts.

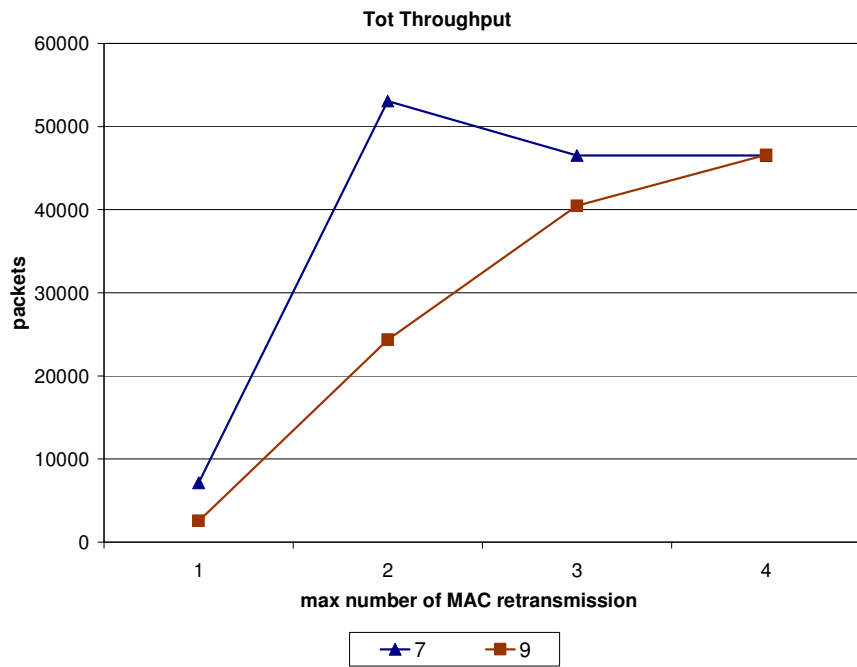


Figure 42. FTP total throughput with different shadowing deviation values; user-AP distance = 10m, MAC queue size = 50pkts.

6 Experimental Results

We present here the most relevant results from the extensive set of simulations we have run. In particular, we first demonstrate how concurrent TCP-based traffic can affect the performance of real time applications. We then compare the outcome with those of our proposed solutions.

6.1 FTP Impact on Real-Time Entertainment Applications

We have intentionally started the various application flows one after the other in order to notice the progressive impact of the successively incoming and overlapping traffic on the preexisting ones. In particular, we expect to witness increasing delays and jitter in the arrival time of packets as we augment the traffic level. However, the bandwidth requirement of the first starting applications in our scenario is well below the effectively available capacity of the IEEE802.11g wireless media. We have to wait until the FTP flow takes action, quickly saturating the channel and the queues along the path with its packets, before being able to clearly detect significant variations in the delays and jitter experienced by the various real time flows. This phenomenon is evident in Fig. 39 and Fig. 40 where the mobile device was located at 10m from a standard IEEE802.11g AP.

6.2 Shadowing and Distance Impact on TCP Throughput

The distance between the AP and the mobile device represents an important factor in determining the transmission quality, especially in a partitioned environment as a house. In fact, Fig. 41 shows that, positioning the mobile device at a distance of 5m from the AP, the maximum throughput achievable by the FTP application is already obtained when utilizing only two retransmissions at the MAC layer. The quality of the transmission signal also depends on the impediments that it may encounter along its path between the AP and the mobile device.




Figure 43. Example of TCP congestion window; max MAC retransmissions = 2.




Figure 44. Example of TCP congestion window; max MAC retransmissions = 3.




Figure 45. FTP total throughput with different MAC queue sizes; user-AP distance = 10m, shadowing deviation = 9.

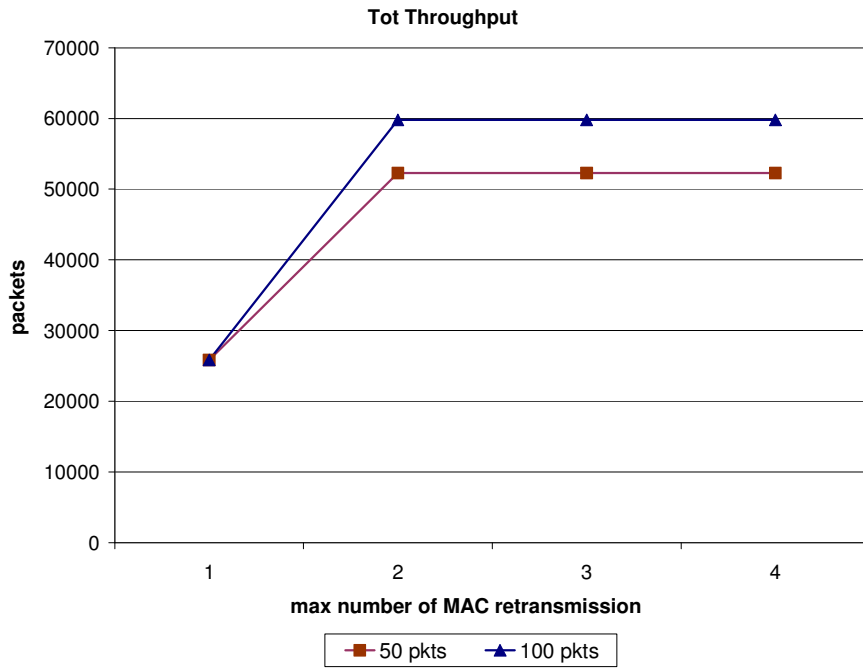


Figure 46. FTP total throughput with different MAC queue sizes; user-AP distance = 5m, shadowing deviation = 9.

We compare now a house environment constituted by medium level partitions (parameter set to 7), with another one having more unfavorable partitions to wireless transmissions (parameter set to 9). Fig. 42 confirms our expectations by showing the throughput gain achievable in the first case. However, Fig. 42 also reports a case where having just two retransmissions at the MAC layer produces a higher total throughput than utilizing a greater number.

TABLE IX. GAMING FLOW JITTER STATISTICS; MAX MAC RETRANSMISSIONS = 4, SHADOWING DEVIATION = 9.

CONSIDERED PERIOD = [0 – 180]s		
Jitter	50 pkts	100 pkts
maximum (ms)	33.740	108.36
average (ms)	1.306	2.041
variance	7.360	22.079
pkts received	2658	2658
CONSIDERED PERIOD = [135 – 180]s		
Jitter	50 pkts	100 pkts
maximum (ms)	33.740	108.36
average (ms)	3.056	5.229
variance	16.665	49.470
pkts received	899	899

TABLE X. GAMING FLOW JITTER STATISTICS; MAX MAC RETRANSMISSIONS = 3, SHADOWING DEVIATION = 9.

CONSIDERED PERIOD = [0 – 180]s		
Jitter	50 pkts	100 pkts
maximum (ms)	31.091	44.632
average (ms)	1.045	1.566
variance	4.833	11.034
pkts received	2654	2655
CONSIDERED PERIOD = [135 – 180]s		
Jitter	50 pkts	100 pkts
maximum (ms)	31.091	44.632
average (ms)	2.292	3.835
variance	11.502	24.431
pkts received	896	897

This apparent paradox has instead a rational explanation. First, we should remember that IEEE802.11 has a limited buffer for transmissions and retransmissions. Second, the game and the video-chat applications generate some reverse traffic that shares path and networking resources with the acknowledgment (ack) packets generated by TCP as transport protocol under the FTP application. When the channel is fully utilized, some acks get lost and may cause timeouts.

As a confirmation of this, Figure 43 and Figure 44 refer to the same configuration of Fig. 42 and present the congestion window, the slow start threshold and the bandwidth delay product of the underlying TCP flow when the maximum number of retransmissions at the MAC layer was set equal to 2 and 3, respectively. The two charts show a higher frequency of timeouts when a higher number of retransmissions at the MAC layer helps the traffic to reach a higher level.

6.3 Appropriately Setting MAC Layer Parameters

Finally we intend to highlight the impact of having different queue sizes and maximum number of retransmissions at the MAC layer on the performance of the various types of traffic present in our considered scenario. Starting with the first parameter, Fig. 45 and Fig. 46 confirm that having larger queue size helps TCP in achieving higher throughputs. However, there is no difference in the achieved throughput when wireless losses, not recovered via MAC retransmissions, are frequent enough to keep the TCP transmission rates low and hence never have the possibility to utilize more than 50 queue slots.

On the other hand, having large queues along the path may augment the total delay time experienced by packets. In fact, each packet waits in queue for a time which proportionally grows with the number of anterior packets already present in the same queue at its arrival. In case of intense traffic, queues tend to be congested and hence queuing delays may become a significant component of the global delays experienced by each packet. At the same time, having larger queue size on a link also spreads the range of possible queuing delays that packets may experience while traveling on that link (depending on the filling level of the queue). The resulting jitter strongly impact on the performance achieved by real time applications and, in particular, by highly interactive applications as video-chats and MMOGs. Statistics of the aforementioned game flow jitter permit a clearer understanding of the performance disparity generated by diverse queue sizes. In particular, the upper part of Table IX refers to the whole simulated duration of the MMOG application, while the rest of it considers only the period when the FTP application is running (from second 135 to second 180).

In any case, the worst jitter is experienced when queues are steadily filled up by the FTP flow. To limit this problem we should find a way to barter part of the FTP throughput with lower queuing delays. This can be attained also by simply reducing the maximum number of retransmission at the MAC layer as it is evident by comparing Table X with Table IX.

Even better jitter average and variance can be gained further diminishing the maximum number of MAC retransmission to 2. However, we advice against this choice, unless placing the device hosting the FTP application closer to the AP or in a house with a better shadowing deviation of the transmitted signal. Otherwise, the FTP throughput descends significantly as can be observed in Fig. 41, Fig. 42, Fig. 45, and Fig. 46.

Summarizing, we can say that a more appropriate configuration of the IEEE802.11g than the traditional one would probably make use of a maximum number of 3 retransmissions, thus guaranteeing a high FTP throughput whilst maintaining a low per-packet delay and jitter. Moreover, when a unique queue is maintained for all the traffic flows, a small size (50 packets) should be preferred.

In particular, for a scenario where the shadowing deviation is set to 7 and we have 10m of distance between the AP and the mobile device, this configuration of the MAC layer parameters allows a TCP goodput of 55371 packets during the 45 seconds when the FTP application was running. The corresponding congestion window, slow start threshold and bandwidth-RTT product is shown in Fig. 47. The interarrival packet time and the jitter for the online game flow traveling from the server to the client are reported in Fig. 48 and Fig. 49, respectively.




Figure 47. TCP congestion window; MAC max retransmissions = 4, buffer size = 50 packets.




Figure 48. Online game interarrival time; MAC max retransmissions = 4, buffer size = 50 packets.




Figure 49. Online game jitter; MAC max retransmissions = 4, buffer size = 50 packets.

6.4 Utilizing TCP Vegas in Place of TCP New Reno

The ability of TCP Vegas in detecting queues and anticipating their further growth is evident in Fig. 50, which shows the congestion window when parameters are set as follows: $\alpha = 1$, $\beta = 3$, $\gamma = 2$. This parameter setting corresponds to utilize a very small amount of buffer at the bottleneck. Consequently, both the queuing time and the achieved goodput will be reduced. The congestion window, in fact, results evidently limited by TCP Vegas algorithm and reaches very low values. The final goodput is obviously affected as well: 43950 packets acknowledged in 45s instead of 46580 and 54137 packets for a buffer size of 50 and 100 packets, respectively, for the regular configuration with TCP New Reno (as seen in Fig. 45).

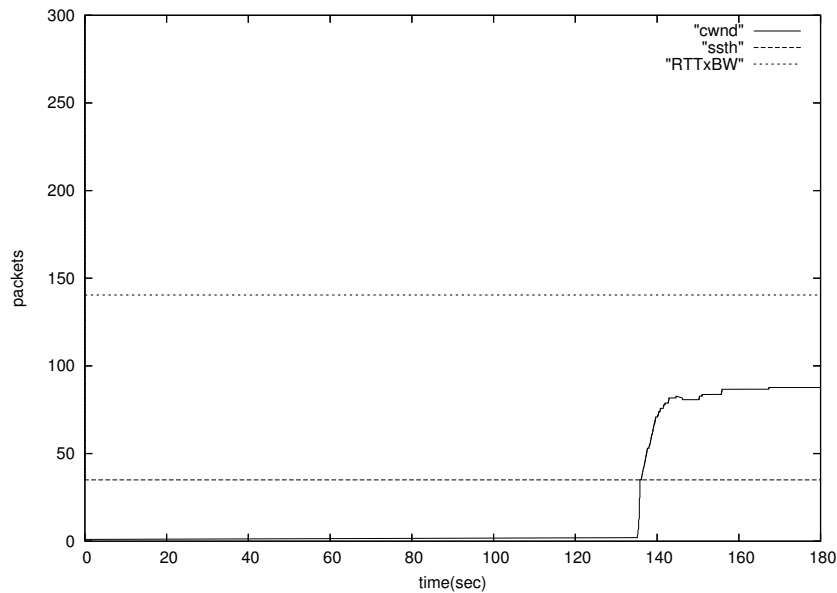


Figure 50. TCP Vegas congestion window; $\alpha = 1$, $\beta = 3$, $\gamma = 2$.




Figure 51. Online game interarrival time with concurrent TCP Vegas; $\alpha = 1$, $\beta = 3$, $\gamma = 2$.




Figure 52. Online game jitter with concurrent TCP Vegas; $\alpha = 1$, $\beta = 3$, $\gamma = 2$.




Figure 53. Online game interarrival time with concurrent TCP Vegas: $\alpha = 5$, $\beta = 10$, $\gamma = 8$.




Figure 54. Online game jitter with concurrent TCP Vegas; $\alpha = 5$, $\beta = 10$, $\gamma = 8$.




Figure 55. TCP Vegas congestion window; $\alpha = 1$, $\beta = 3$, $\gamma = 2$.

On the other hand, having no packets in queue helps real time traffic in reducing its per-packet delay. In particular, in Fig. 51 and Fig. 52 we show the packet interarrival time and the jitter for the online game flow going from the server to the client. The tremendous reduction for these values results evident from the charts: the jitter, for instance, never reaches 4.5 ms.

With different values of its parameters, TCP Vegas becomes able to make use of more buffer space at the AP. This results in higher delays but, at the same time, higher goodputs. Indeed, the tradeoff relationship between the per-packet delay and the total goodput highlighted in Section 5.3 is perfectly embodied in TCP Vegas parameters. By setting the queue target size, those parameters could be seen as knobs able to move that tradeoff towards one direction or the other.

As an example, if we set $\alpha = 5$, $\beta = 10$, and $\gamma = 8$, then the total goodput raises to 59999 packets acknowledged in 45s. On the other hand, the packet interarrival time and the jitter for the online game flow going from the server to the client result increased as shown by Fig. 53 and Fig. 54.

The best tradeoff between goodput and real-time traffic jitter could be obtained for this configuration by using $\alpha = 3$, $\beta = 7$, and $\gamma = 5$. For this set of parameters, in fact, the achieved goodput is still high: 57171. At the same time, the interarrival time and the jitter of real time traffic achieves very good results. Specifically, we can see the trend of the congestion window in Fig. 55, while the interarrival time and the jitter of online game packets going from the server to the client is depicted in Fig. 56 and Fig. 57, respectively.

Unfortunately, TCP Vegas suffers from three main drawbacks which are well known in the scientific community. First, setting its parameters is not a trivial task and depends on many factors such as the buffer size at the bottleneck and the number of flows sharing that link. In particular, the last factor continuously changes and it is not possible to continuously adapt α , β , and γ . Second, TCP Vegas has been shown to be unstable since the contemporary presence of two poles [160]. Third, TCP Vegas behaves very poorly in terms of throughput when competing with the traditional TCP New Reno and TCP Sack (which support the large majority of the data flows in the current Internet) [148]. In presence of congestion, in fact, queues will be fully exploited by traditional TCP, while TCP Vegas will shrink its congestion window as it will sense continuous queue utilization.




Figure 56. Online game interarrival time with concurrent TCP Vegas: $\alpha = 3$, $\beta = 7$, $\gamma = 5$.




Figure 57. Online game jitter with concurrent TCP Vegas; $\alpha = 3$, $\beta = 7$, $\gamma = 5$.

Since these problems and, in particular, the last one, TCP Vegas is not factually deployable in the Internet. We need hence to find an alternative solution that could be practically employed.

6.5 Limiting TCP's Advertised Window

The link capacity actually achieved by a wireless connection is usually less than half of the nominal one. From tests run in our lab with a real 802.11b antenna we were able to get only a maximum of about 5Mbps over a nominal rate of 11Mbps. However this was the result of a tuning of the connectivity obtained by choosing the best positioning for the AP and the mobile device. Under regular circumstances, the factual transmission rate would be even less than that. We obtained coherent results also by simulating an 802.11b link on NS-2.

We have then simulated an 802.11g link on NS-2 and measured a maximum effectively achievable transmission rate of circa 20Mbps. We have enhanced our scenario by enabling the AP to modify the advertised window of returning acks accordingly with the bandwidth left available by the UDP-based flows as determined in (22). In particular, the average UDP-based aggregate traffic was computed through a simple low-pass filter and the new advertised window was determined every 200ms.

In this configuration, various values for the parameter C in (22) have been tested and results are reported in Fig. 58. In this chart, we can see the average, the standard deviation and the maximum value for the delays experienced by the gaming application in its flows directed from the server to the client. Moreover, it also contains the throughput trend of the concurrent TCP flow. We do not report here the equivalent charts corresponding to all the other real time traffic flows since their results are coherent with those presented in Fig. 58 and do not need further explanation.

As clearly shown, both the average and the standard deviation of the online game flow increase when we utilize higher values for C . This is obviously due to the fact that higher C values decrease the resilience to TCP bursts thus leading to some queuing at the AP. However, both the average and the standard deviation are very low for all C values and we could have the wrong impression that the online game flow always experience good performance. Unfortunately, this is not true as can be noticed by looking at the curve representing the maximum delay value experienced by packets traveling through the AP.

Finally, Fig. 58 also demonstrates how the throughput decreases when C is set too low. Instead, if C is set higher than the maximum achievable throughput on the channel (in this case, 20Mbps), then the sender will be allowed to send more packets than those bearable by the bottleneck link causing queuing delays. Moreover, at a certain point, some packets will overflow the buffer and the consequent packet loss will cause the reduction of the sending window and average throughput.

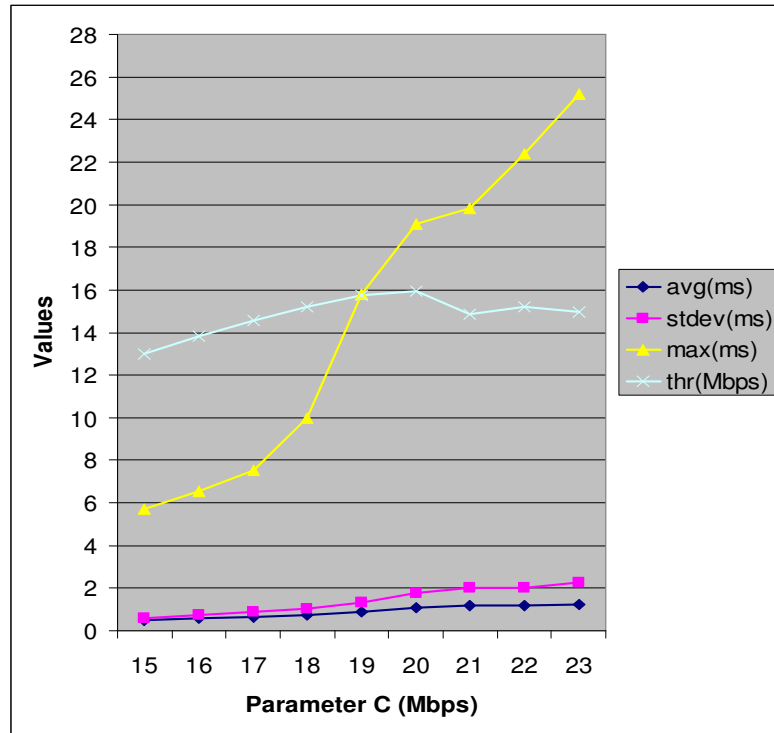


Figure 58. Statistical values when employing limited advertised window.




Figure 59. TCP behavior with limited advertised window and C = 18Mbps.




Figure 60. Online game interarrival time with concurrent TCP, limited advertised window, $C = 18\text{Mbps}$.




Figure 61. Online game jitter with concurrent TCP, limited advertised window, $C = 18\text{Mbps}$.

Setting C as 18Mbps (i.e., the 90% of the maximum achievable bandwidth) seems to be an appropriate choice able to guarantee both low queuing delays and high TCP efficiency. The advertised window exploited by the TCP flow is evident in Fig. 59, which also reports the congestion window, the slow start threshold, and the bandwidth-RTT. We have to keep in mind that the TCP flow starts at second 135 of the simulation time and that the actual sending window is determined as the minimum between the advertised window and the congestion window. Said that, we can appreciate from the chart how the AP is able to keep track of the concurrent real time traffic and determine the most appropriate advertised window. In particular, for this configuration, the final goodput in terms of acknowledged packets over 45 seconds hits 58677, while the interarrival time and the jitter experienced by online game packets are maintained low as demonstrated by Fig. 60 and Fig. 61, respectively.

Another interesting outcome shown by Fig. 59 is the proximity of the advertised window curve and the bandwidth-RTT product ones. The advertised window is prudently set close to the bandwidth-RTT product (i.e., the link pipe size) minus the aggregate UDP based traffic, and this difference also represents an estimate of the amount of real time traffic present on the channel. We have to remind that the real time traffic was simulated considering three simultaneously running applications of various type and exploiting real traffic traces. As it is evident from the chart, the difference between the two curves is relatively small if compared to the whole channel capacity thus demonstrating that real time applications generally do not have to face bandwidth shortage in an 802.11g wireless home, whilst they still have to deal with high and variable delays.

6.6 Summarizing Results

In order to compare the various proposed scheme we have summarized in Fig. 62 statistical results obtained by: i) utilizing regular TCP New Reno on a standard IEEE 802.11g MAC configuration (Regular) ii) appropriately setting the MAC layer parameters (MAC-Setting), iii) utilizing TCP Vegas in place of regular TCP (TCP-Vegas), and iv) TCP with Limited Advertised Window (TCP-LAW) implemented at the AP.

More specifically, the MAC layer parameters for MAC-Setting were set with a maximum number of retransmissions equal to 3 and a buffer size of 50 packets at the AP. Instead, TCP Vegas was configured with $\alpha = 3$, $\beta = 7$, $\gamma = 5$ and TCP-LAW considered $C = 18$.

The compared statistical parameters are the average, the standard deviation and the maximum value of jitter experienced by online game packets entering the house (thus going from the server to the client) via the AP. Results obtained from the other real time applications running in the simulated scenario (i.e. video-stream and video-chat) are coherent with the showed ones; we hence skip to present their charts. Rather, we also show the average throughput achieved by the concurrent TCP connection.

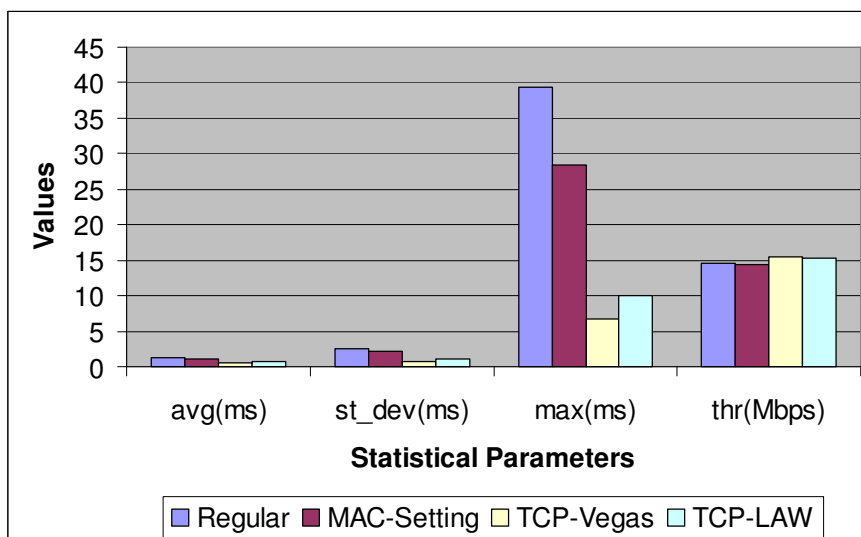


Figure 62. Statistical values for the various schemes.

As it is evident, employing TCP Vegas to support FTP traffic is the solution that would guarantee the best performance both in terms of lowest per packet delay and highest throughput. However, as already mentioned in Section 5.4.2, TCP Vegas cannot be actually deployed in the Internet since it is not able to efficiently coexist with regular TCP flows [148].

Conversely, TCP-LAW could be easily implemented as it only requires the presence of slightly “smarter” APs. The modifications to the AP are very limited thus minimally impacting on their cost and, at the same time, our scheme can perfectly coexist with the current Internet and its employed protocols. Considering this and the remarkable results achieved, TCP-LAW represents the optimal candidate for enhancing wireless home scenarios.

Chapter 6

Conclusion

We are living in a world which is faster and faster spinning toward a continuous, ubiquitous, and seamless wireless connectivity. Customers are more and more attracted by high quality entertainment applications, and mobility represents the next frontier for this kind of technology. On the other hand, real time applications require continuous connectivity and have stringent delays bounds. MMOG represents an exemplar and very interesting case study among the various entertainment applications that would be soon ubiquitously available.

Several factors concur in creating a pleasant game experience for online players: high responsiveness, uniform view of the game state and fairness represent the most important of them. We have hence deeply analyzed issues related to entertainment applications requiring high responsiveness degree as MMOGs.

Particular attention has been devoted to latency issues, since they represent the most important prerequisite in developing online games. We have addressed this problem from two standpoints: the server-to-server synchronization and the client-server communication in a congested wireless environment. We have aimed at creating an effective online game platform able to sustain game sessions even in the challenging scenario involving wirelessly connected players.

In particular, we have first developed a hybrid architecture combining the scalability of peer-to-peer paradigm with the easy management of the client-server one.

We have then devised an efficient synchronization scheme among these servers (GSS) that makes use of a proactive event discarding mechanism, named ILA, to preserve responsiveness. The proposed mechanism relies on the concept of obsolescence and correlation to drop old game events at GSSs, thus accelerating the processing of fresher ones. We have shown simulation results that demonstrate the benefits attainable by employing our scheme in place of traditional ones. In particular, with a limited number of obsolete event drops, ILA is able to proactively guarantee a high interactivity degree also maintaining, in the ILA-RED version, full consistency.

Following a holistic approach we have been able to further refine our synchronization mechanism to improve the degree of fairness experienced by players. We have named this upgraded version FILA and we have presented simulation results showing the improvements in terms of game events that were fairly and interactively delivered to destination with respect to traditional Local Lag schemes.

Finally, to support engaged players even in a mobile scenario, we have proposed and evaluated solutions able to diminish queuing delays even in congested wireless environments. In particular we have devised an easily deployable scheme that exploits enhanced APs and the TCP's advertised window to ensure both goodput efficiency and high responsiveness.

1 Future Work

From here on, our work can proceed toward several directions such as the development of new tools proficient at evaluating MMOG performance or the utilization of SIP technology to support various MMOG functionalities. Among the many possibilities, we elaborate on two potential future studies that

will soon be object of exceptional interest from the scientific community: the IEEE 802.11e standard and the convergence between MMOGs and vehicular networks.

1.1 Utilizing IEEE 802.11e

We have demonstrated how the standard IEEE 802.11 MAC protocol is not well suited for supporting real time traffic and in particular highly interactive applications. Problems arise both from the setting of its parameters, which have been mainly intended for TCP-based traffic, and from the lack of an efficient Quality of Service (QoS) mechanism.

In particular, the regular IEEE 802.11b/g standard implements a contention scheme as previously summarized in Section 5.2. This mechanism is named Distributed Coordination Function (DCF) and supports only best effort services, with no guarantee for bandwidth, packet delay, and jitter. Alternatively, a contention avoidance service intended to better sustain time-bounded traffic makes use of the polling access method offered by the Point Coordination Function (PCF) [161].

PCF scheme provides stations with a prioritized access to the wireless medium. However, several unresolved problems led to the development of enhancements to the protocol which have been very recently finalized in the IEEE 802.11e standard (the definitive document has been released during the final draw up of this Thesis) [102].

Better results in terms of low jitter and per packet delays could be obtained through the utilization of the IEEE 802.11e in place of the regular IEEE 802.11b/g [55, 162]. The IEEE 802.11e defines different classes of traffic and, for each class, diverse configurations of the parameters permit the creation of priorities. These parameters, in fact, determine the various contention delays that packets have to experience before having access to the wireless medium.

From the application point of view, video-chats and online games will belong to the highest priority class of traffic. Since video-stream applications could make use of buffering techniques, they could be featured with a medium priority level. Finally, TCP packets could have a low priority level in order to minimize the impact on real time traffic.

Since the specifications of the IEEE 802.11e have been finally defined only very recently, first experiments will obviously involve simulations. We would like to compare the performance of this new MAC protocol, both in terms of TCP goodput and per-packet delays, with those of the schemes here proposed and, in particular with those of TCP-LAW. Moreover, we would also like to investigate the deployability and the performance of a combination of the two technologies.

1.2 MMOG in Vehicular Networks

Cars represent the next frontier in mobile communication for the forthcoming years, as they become more endowed with entertainment technology. Having 410K cars and SUVs sold in first half of 2004 with pre-installed DVD systems [149], it is not hard to foresee a future where all vehicles will be capable of connecting to the Internet, and passengers will represent a conspicuous slice of the online entertainment market.

In fact, as anticipated in the introduction (Section 1), MMOGs are going to reach also the realm of car-networking. Starting from the assumption of having our mirrored game server architecture endowed with the ILA scheme, we are going to run simulations and real experiments focused on the client-server part of the game platform. In order to test the feasibility of running interactive applications as MMOGs in highly mobile environments, simulations could be run utilizing one of the most popular and trusted simulators in the scientific literature: NS-2 or QualNet [1, 2].

Scenarios involving different kinds of available antennas, as well as diverse dislocation of APs and traffic conditions, could be tested and compared in order to provide useful information about connectivity and latency. Alternative solutions able to compensate AP's lacks of coverage could also exploit car-to-car ad-hoc connections or a combination with 3G full coverage connectivity [11].

In particular, when coverage from APs is missing, some cars could be connected through 3G architecture. These cars could download and upload game packets for themselves but also on behalf of other players sit in cars around them. The general small size of game packets, in fact, limits the bandwidth requirement. Therefore, one single connection could be enough to let several players updating their game status.

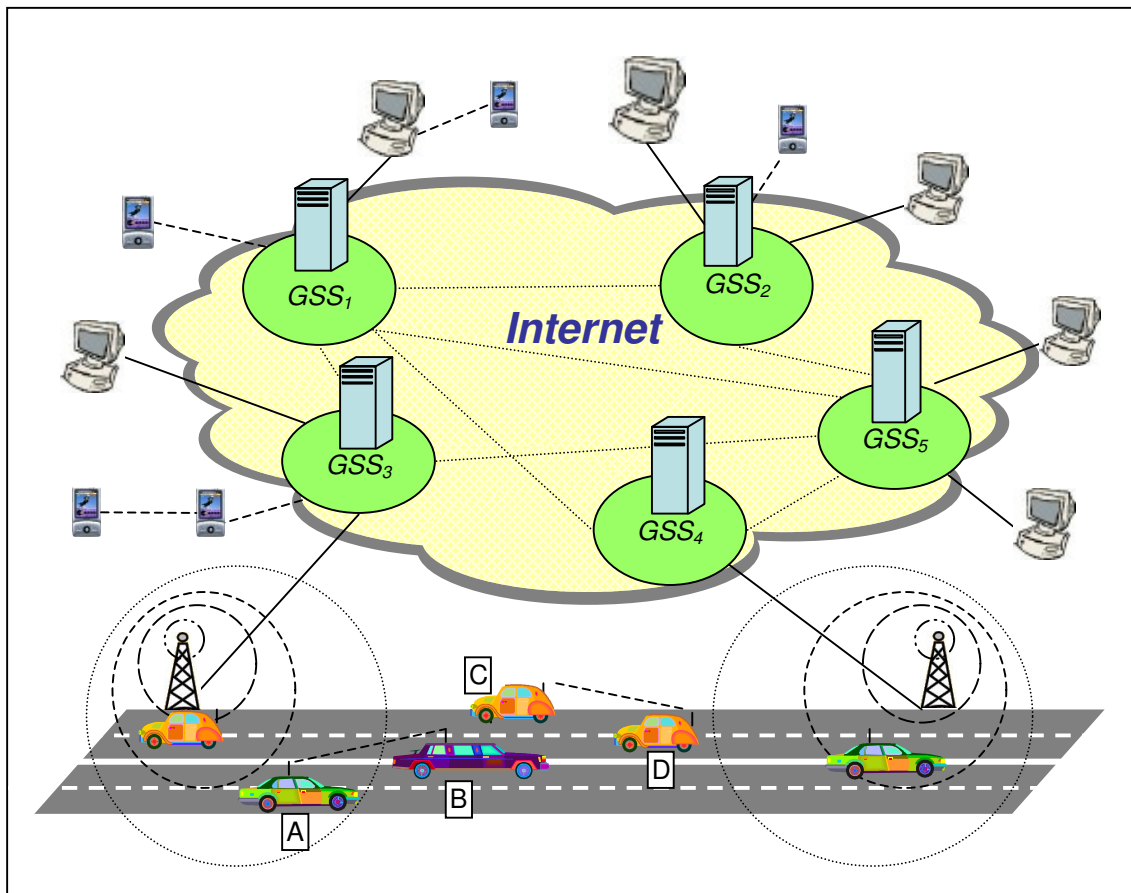


Figure 63. Hybrid architecture for distributed game entertainment in heterogeneous scenarios including car-networking.

This is useful for two main reasons: first, to increase throughput and latency efficiency, and second, to circumvent the current limitation that imposes boundaries on the number of contemporary connected customers for each cell [11]. Wi-Fi connections exploiting car-to-car ad-hoc topology will then be used to diffuse game events between relay vehicles and others in proximity. Cellular and ad-hoc network architecture may thus be both present and combined on each automobile to ensure a continuous connectivity while waiting for being in range of the next AP on the road (as depicted in Fig. 63). Previous studies demonstrated how the aggregate throughput may increase by utilizing a unified cellular and had-hoc network architecture [21]. However, the latency issue involved in this combination has not yet been investigated. Moreover, car-networking scenario further increases the mobility, thus asking for frequent changes of utilized access technology.

Talking about 3G communications, pricing techniques should be investigated in order to put in practice a scenario where people are continuously connected also to propagate other customers' traffic. A mechanism aimed at billing the right final player even when not directly connected to a 3G antenna should be defined. Each player should pay for its own traffic, and only for it. A possible alternative is represented by having flat rates for all the customers. With this solution, players would pay just a monthly fee for their 3G traffic thus increasing the appeal of this architecture and simplifying pricing problems especially in our considered hybrid scenario.

Obviously, simulations will result very helpful at the beginning to evaluate novel solutions in this field. Indeed, real experiments aimed at testing connectivity under diverse conditions could require the deployment of unsustainable investments to provide an adequate amount of participating cars and antennas on the road.

However, simple but meaningful real experiments can be performed to obtain real measurements about the time involved in connectivity settling, handovers, car-to-car communication and different access technology adopted. A limited number of moving automobiles can be used to carry out meaningful experiments. Equipped with devices to communicate with APs, these vehicles will transport users engaged in MMOGs. Even simpler but still valuable experiments can be run inside a laboratory, where a couple of APs and some wireless connected devices could permit some useful delay and connectivity measurements.

Acknowledgment

This PhD thesis has been developed as part of my joint PhD program. Thanks to the Interlink project, in fact, I am a PhD student in both the Dipartimento di Scienze dell'Informazione, Università di Bologna, and in the Department of Computer Science, UCLA. As part of my education and research work, every year I spend a period in each of these two institutions, under the supervision of Prof. Marco Rocchetti and Prof. Mario Gerla, respectively. My first thanks goes hence to them for the great opportunity they gave me and for their fundamental guidance in these years.

I cannot forget Giovanni Pau that conceived, first, this great adventure of a joint PhD. He has been a true friend, a great co-worker, and an unlimited source of astute suggestions.

Many thanks also to Stefano Cacciaguerra, Maria Fazio, Stefano Ferretti, Gustavo Marfia, Daniela Maniezzo, Matteo Roffilli, and Cesare Roseti. With them my papers have been better works, my time in the office enjoyable, and my life richer of precious friends. My thanks have to be extended also to all the people in the Network Research Lab, Radio Lab, and Studio 16 for all the interesting discussion we had.

Finally I wish to thank the reviewers of this thesis for their important comments and suggestions.

This work is dedicated with love to my family.

References

- [1] The Network Simulator, NS-2. <http://www.isi.edu/nsnam/ns/>
- [2] QualNet Network Simulator. <http://www.qualnet.com/>
- [3] C.H. Nam, Soung C. Liew, C.P. Fu, "An Experimental Study of ARQ Protocol in 802.11b Wireless LAN", in Proc. of IEEE VTC 2003, 2003.
- [4] G. Xylomenos and G. C. Polyzos, "TCP and UDP Performance over a Wireless LAN", in Proc. of IEEE INFOCOM '99, 1999.
- [5] C. E. Palazzi, C. Roseti, M. Luglio, M. Gerla, M. Y. Sanadidi, and J. Stepanek, "Satellite coverage in urban areas using Unmanned Airborne Vehicles (UAVs)", in Proc. of VTC2004 Spring, Milan, Italy, May 2004.
- [6] C. E. Palazzi, C. Roseti, M. Luglio, M. Gerla, M. Y. Sanadidi, and J. Stepanek, "Enhancing Transport Layer Capability in HAPS-Satellite Integrated Architecture", Wireless Personal Communications International Journal, Special Issue on "High Altitude Platforms: Research and Application Activities", Kluwer Academic Publishers, 2004.
- [7] C. Griwodz, "State Replication for Multiplayer Games", in Proc. of NetGames2002, pp. 29-35, Braunschweig, Germany, 2002.
- [8] S. Mascolo, A. Grieco, G. Pau, M. Gerla, C. Casetti, "End-to-End Bandwidth Estimation in TCP to Improve Wireless Link Utilization", European Wireless Conference, Florence, Italy, Feb 2002.
- [9] M. Rocchetti, V. Ghini and P. Salomoni, "Distributing Music from IP Networks to UMTS Terminals: an Experimental Study", in Proc. 2002 SCS Euromedia Conference (EUROMEDIA2002), (M. Rocchetti Ed.) The Society for Modeling and Simulation International, Modena (Italy), pp. 147-154, Apr 2002.
- [10] V. Ghini, G. Pau, M. Rocchetti, P. Salomoni, M. Gerla, "For Here or To Go? Downloading Music on the Move with an Ultra Reliable Wireless Internet Application", IEEE ICC'2004, Paris, 2004.
- [11] UMTS Forum. <http://www.umts-forum.org>
- [12] J. Weatherall and A. Jones, "Ubiquitous Networks and their applications", in IEEE Wireless Communication, Feb 2002.
- [13] E. Gustafsson and A. Jonsson, "Always Best Connected", IEEE Wireless Communications, pp. 49-55, Feb 2003.
- [14] G. Huston, "The future for TCP", The Internet Protocol Journal, vol. 3, n. 3, Sep 2000.
- [15] G. Huston, "TCP in a Wireless World", IEEE Internet Computing, pp. 82 – 84, Mar-Apr 2001.
- [16] Mattias Esbjörnsson, Oskar Juhlin and Mattias Östergren, "The Hocman Prototype – Fast Motor Bidders and Ad-hoc Networking", in Proc. of MUM, 2002.
- [17] H. Balakrishnan, V. N. Padmanabhan, S. Sehan, and R. H. Katz, "A comparison of mechanism for improving TCP performance over wireless links", IEEE/ACM Trans. Networking, vol. 5, no. 6, pp. 756 – 769, Dec 1997.
- [18] M. Mauve, J. Widmer and H. Hartenstein, "A Survey on Position-Based Routing in Mobile Ad Hoc Networks", IEEE Network, vol. 15, no. 6., pp. 30-39, 2001.
- [19] J. Ott and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for Automobile Users", In Proc. of IEEE INFOCOM 2004, Mar 2004.

- [20] A. Nandan, S. Das, G. Pau., M. Gerla, M. Y. Sanadidi, "Co-operative Downloading in Vehicular Ad-hoc Wireless Networks", in Proc. of WONS 2005, Saint Moritz, Switzerland, Jan 2005.
- [21] H. Luo, R. Ramjee, P. Sinha, L. Li, S. Lu, "UCAN: A Unified Cellular and Ad-Hoc Network Architecture", in Proc. of ACM MobiCom 2003, San Diego, CA, Sep 2003.
- [22] Q. Xu, T. Mak, J. Ko, R. Sengupta, "Vehicle-to-Vehicle Safety Messaging in DSRC", in Proc. of ACM VANET'04, Philadelphia, PA, Oct 2004.
- [23] B. Hughes, R. Meier, R. Cunningham, V. Cahill, "Towards Real-Time Middleware for Vehicular Ad Hoc Networks", in Proc. of ACM VANET'04, Philadelphia, PA, Oct 2004.
- [24] Mobile Monday. http://www.mobilemonday.net/mm/story.php?story_id=3748
- [25] MMOGCHART. <http://www.mmogchart.com/>
- [26] D. Kushner, Masters of Doom, Random House, New York, NY, 2002.
- [27] J. Smed, T. Kaukoranta, H. Hakonen, "Aspects of Networking in Multiplayer Computer Games", in Proc. of International Conference on Application and Development of Computer Games in the 21st Century, pp.74-81, Hong Kong, China, 2001.
- [28] M. R. Macedonia, A Network Software Architecture for Large Scale Virtual Environments, Ph.D. Thesis, Naval Postgraduate School, Monterey, CA, 1995.
- [29] D. L. Neyland Virtual Combat: A Guide to Distributed Interactive Simulation, Stackpole Books, Mechanicsburg, PA, 1997.
- [30] V. Normand, "The COVEN project: Exploring Applicative, Technical, and Usage Dimensions of Collaborative Virtual Environments", Presence, vol.8, no.2, pp. 218-236, 1999.
- [31] E. Frécon, M. Stenius, "DIVE: A Scalable Network Architecture for Distributed Virtual Environments", Distributed Systems Engineering, vol. 5, no. 3, pp. 91-100, 1998.
- [32] T. A. Funkhouser, "RING: A Client-Server System for Multi-User Virtual Environments", in Proc. of the 1995 Symposium on Interactive 3D Graphics, pp. 85-92, Monterey, CA, 1995.
- [33] Nintendo DS. <http://www.nintendo.com/ds/index.jsp>
- [34] Sony PSP. <http://www.us.playstation.com/pressreleases.aspx?id=207>
- [35] Quake Forge Project. <http://www.quakeforge.org/>
- [36] R. M. Mine, J. Shochet, R. Hughston, "Building a Massively Multiplayer Game for the Million: Disney's Toontown Online", ACM Computers in Entertainment (CIE), vol. 1, no. 1, pp. 15-15, 2003.
- [37] P. Mehta and S. Udani. "Voice over IP", in IEEE Potentials, vol. 20, no. 4, Oct 2001.
- [38] W. Cai, P. Xavier, S. J. Turner, B. Lee, "A Scalable Architecture for Supporting Interactive Games on the Internet", in Proc. of the 16th Workshop on Parallel and Distributed Simulation, pp. 54-61, Washington, DC, 2002.
- [39] S. Singhal, M. Zyda, Networked Virtual Environments: Design and Implementation, Addison Wesley, 1999.
- [40] T. Kanter and C. Olrog, "VoIP in Applications for Wireless Access", in IEEE Workshop on Local and Metropolitan Area Networks, Nov 1999.
- [41] K. Svanbro, J. Wiorek, and B. Olin, "Voice-over-IP-over-wireless", in IEEE PIMRC '00, Sep 2000.

- [42] C. E. Palazzi, "Buddy-Finder: A Proposal for a Novel Entertainment Application for GSM", in Proc. of the 1st IEEE International Workshop on Networking Issues in Multimedia Entertainment (NIME'04), GLOBECOM 2004, Dallas, TX, Nov 2004.
- [43] C. E. Palazzi, "Residual Capacity Estimator for TCP on wired/wireless links", in Proc. of WCC2004-Student-Forum IFIP World Computer Congress 2004, Toulouse, France, Aug 2004.
- [44] S. Tekinay and B. Jabbari, "Handover and channel assignment in mobile cellular networks", in IEEE Communication Magazine, vol. 29, no, 11, pp. 42-46, Nov 1991.
- [45] N. Davies, K. Cheverst, A. Friday and K. Mitchell, "Future Wireless Applications for a Networked City: Services for Visitors and Residents", Feb 2002.
- [46] V. Ghini, G. Pau, M. Roccetti, P. Salomoni and M. Gerla, "Smart Download on the Go: A Wireless Internet Application for Music Distribution over heterogeneous Networks", in Proc. IEEE International Conference on Communications - Access and Home Networks Symposium (ICC'04), accepted for publication, Paris, France, IEEE Communications Society, Jun 2004.
- [47] M. Roccetti and P. Salomoni, "The Design and Performance of a Wireless Internet Application for Supporting Multimedia City Guides", in Proc. of IEEE International Conference on Information Technology: Research and Education (ITRE 2003), Special Session on Multimedia Transport in Heterogeneous Wireless Networks, Newark (USA), Aug 2003.
- [48] M. Stemm and R. H. Katz. "Vertical handoffs in wireless overlay networks", in Mobile Networks and Applications, vol. 3, no. 4, 1998.
- [49] K. Wang, S.K. Tripathi, "Mobile-End Transport Protocols: An Alternative to TCP/IP Over Wireless Links", in Proc. IEEE Infocom '98, San Francisco, California, USA, pp. 1046-1053, Apr 1998.
- [50] M. Stemm and R. H. Katz, "Vertical Handoffs in Wireless Overlay Networks", ACM Mobile Networks and Applications (MONET) Journal, Special Issue on Mobile Networking in the Internet, 1998.
- [51] K. Brown, S. Singh, "M-TCP: TCP for Mobile Cellular Networks", ACM Computer Communication Review, vol. 27, no. 5, Jul 1997.
- [52] J. S. Steinman, R. Bagrodia, D. Jefferson, "Breathing Time Warp", in Proc. of the 1993 Workshop on Parallel and Distributed Simulation, San Diego, CA pp. 109-118, 1993.
- [53] D. R. Jefferson, "Virtual Time", ACM Transaction on Programming Languages and Systems, vol. 7, no. 3, pp. 404-425, 1985.
- [54] S. J. Kim, F. Kuester, K. H. Kim "A Global Timestamp-based Scalable Framework for Multi-player Online Games", in Proc. of the IEEE Fourth International Symposium on Multimedia Software Engineering (MSE'02), 2002.
- [55] H. Yoon, J. W. Kim, D. Y. Shin, "Dynamic Admission Control in IEEE 802.11e EDCA-based Wireless Home Network", in Proc. of IEEE Consumer Communications and Networking Conference (CCNC 2006), Las Vegas, NV, Jan 2006.
- [56] M. Gerla, K. Tang, R. Bagrodia, "TCP Performance in Wireless Multi-hop Networks", in Proc. of the 2nd IEEE Workshop on Mobile Computer Systems and Applications (WMCSA '99). Washington, DC, USA, p. 41-50, 1999.
- [57] M. Heusse, F. Rousseau, G. Berger-Sabbatel, A. Duda, "Performance Anomaly of 802.11b", in Proc. of IEEE INFOCOM 2003, San Francisco, CA, USA, Apr 2003.
- [58] T. Beigbeder, R. Coughlan, C. Lusher, J. Plunkett, E. Agu, M. Claypool, "The Effects of Loss and Latency on User Performance in Unreal Tournament 2003", in Proc. of ACM Network and System Support for Games Workshop (NetGames), Portland, OG, USA, Sep 2004.

- [59] E. Cronin, A. R. Kurc, B. Filstrup, S. Jamin, "An Efficient Synchronization Mechanism for Mirrored Game Architectures", *Multimedia Tools and Applications*, vol. 23, no. 1, pp. 7-30, 2004.
- [60] D. R. Cheriton, D. Skeen, "Understanding the Limitations of Causal and Totally Ordered Multicast", in *Proc. of the 14th Symposium on Operating System Principles (SOSP '93)*, pp. 44-57, Asheville, NC, 1993.
- [61] X. Défago, A. Schiper, P. Urban, "Totally Ordered Broadcast and Multicast Algorithms: a Comprehensive Study", Technical Report, DSC/2000/036, Swiss Federal Ecole Polytechnique Fédérale de Lausanne, Switzerland, 2000.
- [62] C. E. Palazzi, S. Ferretti, S. Cacciaguerra, M. Roccetti, "On Maintaining Interactivity in Event Delivery Synchronization for Mirrored Game Architectures", in *Proc. of 1st IEEE International Workshop on Networking Issues in Multimedia Entertainment (NIME'04), GLOBECOM 2004*, Dallas, TX, 2004.
- [63] S. Ferretti, M. Roccetti, "A Novel Obsolescence-based approach to Event Delivery Synchronization in Multiplayer Games", *International Journal of Intelligent Games and Simulation*, vol. 3, no. 1, pp. 7-19, 2004.
- [64] S. Ferretti, M. Roccetti, S. Cacciaguerra, "On Distributing Interactive Storytelling: Issues of Event Synchronization and a Solution", in *Proc. of the 2nd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2004)*, LNCS 3105, Darmstadt, Germany, pp. 219-231, 2004.
- [65] R. Drummond, O. Babaoglu, "Low-Cost Clock Synchronization", *Distributed Computing*, vol. 6, no. 3, pp. 193-203, 1993.
- [66] F. Cristian, "Probabilistic clock synchronization", *Distributed Computing*, vol. 3, no. 3, pp. 146-158, 1989.
- [67] D. L. Mills, "Internet Time Synchronization: the Network Time Protocol", *IEEE Transactions on Communications*, vol. 39, no. 10 pp. 1482-1493, 1991.
- [68] S. Floyd, V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance", *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397-413, 1993.
- [69] C. E. Palazzi, S. Ferretti, S. Cacciaguerra, M. Roccetti, "A RIO-like Technique for Interactivity Loss Avoidance in Fast-Paced Multiplayer Online Games: a Preliminary Study", in *Proc. of the 2nd Annual International Workshop in Computer Game Design and Technology (GDTW 2004)*, Liverpool, UK, 2004.
- [70] D. D. Clark, W. Fang, "Explicit Allocation of Best-Effort Packet Delivery Service", *IEEE/ACM Transactions on Networking*, vol.6, no.4, pp.362-373, 1998.
- [71] M. Mauve, *Distributed Interactive Media*, PhD thesis, University of Mannheim, ISBN 3-89838-471-3. infix., Berlin, 2000.
- [72] K. Park, W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*, Wiley-Interscience, 1st Edition, 2000.
- [73] J. Farber, "Network Game Traffic Modelling" in *Proc. of NetGames2002*, pp.53-57, Braunschweig, Germany, 2002.
- [74] Half-Life. <http://counterstrike.sierra.com/>
- [75] S. Wright, S. Tischer, "Architectural Considerations in Online Game Services over DSL Networks", in *Proc. IEEE International Conference on Communications - (ICC'04)*, IEEE Communications Society, Paris, France, 2004.

- [76] G. Armitage, "An Experimental Estimation of Latency Sensitivity in Multiplayer Quake 3", in Proc. of ICON, pp. 137-141, Sydney, Australia, 2003.
- [77] M. S. Borella, "Source Models for Network Game Traffic", Computer Communications, vol. 23, no. 4, pp. 403-410, 2000.
- [78] L. Pantel, L. C. Wolf, "On the Impact of Delay on Real-Time Multiplayer Games", in Proc of the 12th International Workshop on Network and Operating Systems Support for Digital Audio and Video, Miami, FL, May 2002.
- [79] T. Henderson, "Latency and User Behaviour on a Multiplayer Game Server", in Proc. of the 3rd International Workshop on Networked Group Communication (NGC01) pp. 1-13, London, UK, Nov 2001.
- [80] F. Fitzek, G. Schulte, M. Reisslein, "System Architecture for Billing of Multi-Player Games in a Wireless Environment Using GSM/UMTS and WLAN Services", in Proc. of NetGames2002, pp. 58-64, Bruanschweig, Germany, Apr 2002.
- [81] C. E. Palazzi, S. Ferretti, S. Cacciaguerra, M. Rocchetti, "Interactivity-Loss Avoidance in Event Delivery Synchronization for Mirrored Game Architectures", IEEE Transactions on Multimedia, Oct 2005. *Accepted for publication.*
- [82] Mesh Meets Ad Hoc: the Urban Vehicular Grid. <http://www.cs.ucla.edu/ST/>
- [83] Methods for Subjective Determination of Transmission Quality, ITU-T Recommendation P.800, Aug 1996.
- [84] A. Rix, R. Reynolds, M. Hollier, "Perceptual Measurement of End-to-End Speech Quality over Audio and Packet-Based Networks", AES 106th Convention, Munich, Germany, May 1999.
- [85] Methods for the Subjective Assessment of Small Impairments in Audio Systems Including Multichannel Sound Systems, ITU-R Recommendation BS.1116, Jul 1998.
- [86] Method for Objective Measurements of Perceived Audio Quality, ITU-R Recommendation BS.1387, Jan 1999.
- [87] A. Rix, R. Reynolds, M. Hollier, "Robust Perceptual Assessment of End-to-End Audio Quality", in Proc. of 1999 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, New Paltz, NY, Oct 1999.
- [88] M. Handley et al., "SIP: Session Initiation Protocol", IETF RFC 2543, Mar 1999.
- [89] H. Hartenstein, B. Bochow, A. Ebner, M. Lott, M. Radimirsch, D. Vollmer, "Position-Aware Ad Hoc Networks for Inter-vehicle Communications: The Fleenet Project", in ACM Symposium on Mobile Ad Hoc Networking and Computing, MobiHOC, Oct 2001.
- [90] Federal Communications Commission, FCC 99-305, FCC Report and Order, Oct 1999.
- [91] Federal Communications Commission, FCC 03-024, FCC Report and Order, Feb 2004.
- [92] J. Jun and M. L. Sichitiu, "The Nominal Capacity of Wireless Mesh Networks", IEEE Wireless Communications magazine, Oct 2003.
- [93] I. Foster and C. Kesselman, editors. The Grid: Blueprint for a New Computing Infrastructure. Morgan-Kaufmann, 1999.
- [94] L. Gautier, C. Diot, "Design and Evaluation of MiMaze, a Multi-player Game on the Internet", in Proc. of IEEE Multimedia (ICMCS'98), pp. 234-236 Austin, TX, USA, 1998.
- [95] Everquest. <http://www.everquest.com>
- [96] Ultima Online. <http://www.uo.com>

- [97] Butterfly Grid Solution for Online Games. <http://www.butterfly.net>
- [98] J. S. Steinman, "Scalable Parallel and Distributed Military Simulations Using the SPEEDES Framework", in Proc. of 2nd Electronic Simulation Conference (ELECSIM95), Internet, 1995.
- [99] A. Balk, M. Gerla, M. Sanadidi, D. Maggiorini "Adaptive Video Streaming: Pre-encoded MPEG-4 with Bandwidth Scaling", Computer Networks vol. 44, no. 4, pp. 415-439, 2004.
- [100] C. Boyd, A. Mathuria, "Key Establishment Protocols for Secure Mobile Communications: a Critical Survey", Computer Communications, vol. 23, pp. 575-587, 2000.
- [101] P. Mohapatra, J. Li, C. Gui, "QoS in Mobile Ad Hoc Networks", IEEE Wireless Communications magazine, Jun 2003.
- [102] IEEE Standard for Information Technology, "Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks - Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment: Medium Access Control (MAC) Quality of Service Enhancements", P802.11e/D13.0, 2005.
- [103] S. Cacciaguerra, M. Rocchetti, M. Roffilli, A. Lomi, "A Wireless Software Architecture for Fast 3D Rendering of Agent-Based Multimedia Simulations on Portable Devices" in Proc. of the First Consumer Communications and Networking Conference (CCNC), IEEE Communications Society, Las Vegas, NV, Jan 2004.
- [104] S. Zander, I. Leeder, G. Armitage, "Achieving Fairness in Multiplayer Network Games through Automated Latency Balancing", in Proc. of ACM SIGCHI ACE2005, pp. 117-124, Valencia, Spain, 2005.
- [105] M. Mauve, J. Vogel, V. Hilt, W. Effelsberg, "Local-lag and Timewarp: Providing Consistency for Replicated Continuous Applications", IEEE Transactions on Multimedia, vol. 6, no. 1, pp. 47-57, 2004.
- [106] F. Fitzek, G. Schulte, M. Reisslein, "System Architecture for Billing of Multi-Player Games in a Wireless Environment Using GSM/UMTS and WLAN Services", in Proc. of the 1st Workshop on Network and System Support for Games (NetGames2002), ACM SIG MM, pp. 58-64, Braunschweig, Germany, 2002.
- [108] Sheldon N., Girard E., Borg S., Claypool M., Agu E. The Effect of Latency on User Performance in Warcraft III. In Proceedings of the 2nd Workshop on Network and System Support for Games (NetGames 2003), ACM SIGCOMM and SIG MM, pp. 3-14, Redwood City, CA, USA, 2003.
- [109] S. Ferretti, C. E. Palazzi, M. Rocchetti, M. Gerla, G. Pau, "*Buscar el Levante por el Poniente*: In Search of Fairness Through Interactivity in Massively Multiplayer Online Games", in Proc. of the 2nd IEEE International Workshop on Networking Issues in Multimedia Entertainment (NIME'06), CCNC 2006, Las Vegas, NV, Jan 2006.
- [110] T. Jehaes, D. De Vleeschauwer, T. Coppens, B. Van Doorselaer, E. Deckers, W. Naudts, K. Spruyt, R. Smets, "Access Network Delay in Networked Games", in Proc. of the 2nd Workshop on Network and System Support for Games (NetGames 2003), ACM SIGCOMM and SIG MM, pp. 63-71, Redwood City, CA, USA, 2003.
- [111] P. Ramanathan, K. G. Shin, R. W. Butler, "Fault Tolerant Clock Synchronization in Distributed Systems", IEEE Computer, vol. 23, no. 10, pp. 33-42, 1990.
- [112] F. Safaei, P. Boustead, C. D. Nguyen, J. Brun, M. Dowlatshahi, "Latency Driven Distribution: Infrastructure Needs of Participatory Entertainment Applications", IEEE Communications Magazine, IEEE Communications Society, Special Issue on "Entertainment Everywhere: System and Networking Issues in Emerging Network-Centric Entertainments Systems", Part I, May 2005.
- [113] Skype for Pocket PC. <http://www.skype.com/products/skype/pocketpc/>

- [114] A. Nandan, S. Das, B. Zhou, G. Pau, M. Gerla, "AdTorrent: Digital Billboards for Vehicular Networks", in Proc. of IEEE/ACM International Workshop on Vehicle-to-Vehicle Communications (V2VCOM), San Diego, CA, USA, Jul 2005.
- [115] A. Nandan, S. Das, G. Pau, M.Y. Sanadidi, M. Gerla "Cooperative Downloading in Vehicular Ad Hoc Wireless Networks", in Proc. of IEEE/IFIP International Conference on Wireless On demand Network Systems and Services, St. Moritz, Switzerland, pp. 32-41, Jan 2005.
- [116] S. Das, A. Nandan, G. Pau, M.Y. Sanadidi, M. Gerla "SPAWN: Swarming Protocols for Vehicular Ad Hoc Wireless Networks", in Proc. of the First ACM International Workshop on Vehicular Ad Hoc Networks (VANET 2004), MOBICOM 2004, Berkeley, CA, USA, pp. 93-94, 2004.
- [117] S. Moyer, D. Marples, S. Tsang, and A. Ghosh, "Service Portability of Networked Appliances," 2002. [Online]. Available at: <http://citeseer.ist.psu.edu/moyer00service.html>
- [118] A. Mani, H. Sundaram, D. Birchfield, and G. Qian, "The Networked Home as a User-Centric Multimedia System," in NRBC '04: Proceedings of the 2004 ACM Workshop on Next-Generation Residential Broadband Challenges. New York, NY, USA: ACM Press, pp. 19-30, 2004.
- [119] G. Bell, J. Gemmell, "A Call for the Home Media Network", Communications of the ACM, vol. 45, no. 7, pp. 71-75, 2002.
- [120] M. Tsang, G. Fitzmaurice, G. Kurtenbach, A. Khan, "Game-Like Navigation and Responsiveness in Non-Game Applications", Communications of the ACM, vol. 46, no. 7, pp. 56-61, Jun 2003.
- [121] A. E. Cha, "Recalling Iraq's Terrors Through Virtual Reality", Washington Post, Mar 2005. <http://www.washingtonpost.com/wp-dyn/articles/A58360-2005Mar22.html>
- [122] A. Spagnolli, L. Gamberini, F. Scarpetta, S. Colognesi, "Ironia e ristrutturazione del participant framework: il caso degli ambienti virtuali", Rivista di Psicolinguistica Applicata, vol. 2, pp. 165-176, 2004.
- [123] A. Spagnolli, L. Gamberini, P. Cottone, G. Mantovani, "Ergonomics of Virtual Environments for Clinical Applications", Internet and Virtual Reality as Assessment and Rehabilitation Tools for Clinical Psychology and Neuroscience, Amsterdam, IOS Press (Netherlands), pp. 217-230, 2004.
- [124] M. Billinghurst, H Kato, "Collaborative Augments Reality", Communications of the ACM, vol. 45, no. 7, pp. 64-70, Jul 2002.
- [125] S. Ferretti, M. Rocchetti, "Event Synchronization for Interactive Cyberdrama Generation on the Web: a Distributed Approach", in Proc. of 13th International World Wide Web Conference (WWW 2004), vol. WWW2004 Poster Track, W3C/ACM editor, New York, NY, USA, May 2004.
- [126] T. M. Rhyne, "Computer Games and Scientific Visualization", Communications of the ACM, vol. 45, no. 7, pp. 40-44, Jul 2002.
- [127] J. Brun, P. Boustead, F. Safaei, "Fairness and Playability in Online Multiplayer Games", in Proc. of the 2nd IEEE International Workshop on Networking Issues in Multimedia Entertainment (NIME'06), CCNC 2006, Las Vegas, NV, Jan 2006.
- [128] C. Wagner, M. Schill, R. Manner, "Intraocular Surgery on a Virtual Eye", Communications of the ACM, vol. 45, no. 7 pp. 45-49, Jul 2002.
- [129] J. Kumagai, "Fighting in the streets", IEEE Spectrum, vol. 38, no. 2, pp. 68-71, Feb 2001.
- [130] The Eight Fallacies of Distributed Computing. <http://today.java.net/jag/Fallacies.html>

- [131] M. Mauve, S. Fischer, J. Widmer, "A Generic Proxy System for Networked Computer Games", in Proc. of the 1st Workshop on Network and System Support for Games, pp. 25-28, 2002.
- [132] C. Diot, L. Gautier, "A Distributed Architecture for Multiplayer Interactive Applications on the Internet", IEEE Network Magazine, vol. 13, no. 4, Jul/Aug 1999.
- [133] K. W. Lee, B. J. Ko, S. Calo, "Adaptive Server Selection for Large Scale Interactive Online Games", in Proc. of the 14th International Workshop on Network and Operating Systems Support for Digital Audio and Video, pp. 152-157, 2004.
- [134] K. Alex, S. Taylor, "Using Determinism to Improve the Accuracy of Dead Reckoning Algorithms", in Proc. of Simulation Technologies and Training Conference, Sydney, 2000.
- [135] L. Pantel, L. C. Wolf, "On the Suitability of Dead Reckoning Schemes for Games", in Proc. of the 1st Workshop on Network and System Support for Games, pp. 79-84, 2002.
- [136] B. Knutsson, H. Lu, W. Xu, B. Hopkins, "Peer-to-Peer Support for Massively Multiplayer Games", in Proc. of the Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2004), pp. 96-107, Mar 2004.
- [137] F. W. B. Li, L. W. F. Li, R. W. H. Lau, "Supporting Continuous Consistency in Multiplayer Online Games", in Proc. of the 12th Annual ACM International Conference on Multimedia, pp. 388-391, 2004.
- [138] J. Postel, "Rfc0793: Transmission Control Protocol," Internet Engineering Task Force (IETF), Tech. Rep., 1981.
- [139] V. Jacobson, "Congestion Avoidance and Control", in Proc. of ACM SIGCOMM'88, pp. 314-329, 1988.
- [140] S. Floyd, "A Report on Recent Developments in TCP Congestion Control", IEEE Communications, vol. 39, no. 4, pp. 84-90, 2001.
- [141] Y. M. Wang, W. Russell, A. Arora, R. K., Jagannathan, J. Xu, "Towards Dependable Home Networking: an Experience Report", in Proc. of IEEE International Conference on Dependable Systems and Networks (DSN 2000), New York, USA, pp. 43-49, Jun 2000.
- [142] IEEE, "Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (Phy) Specifications," Specifications, ISO/IEC 8802-11:1999(E), 1999.
- [143] S. Low, L. Peterson, L. Wang, "Understanding Vegas: a Duality Model," Journal of ACM, vol. 49, no. 2, pp. 207-235, Mar 2002.
- [144] L. Brakmo, S. O'Malley, L. Peterson, "TCP Vegas: New Techniques for Congestion Detection and Avoidance", in Proc. of the SIGCOMM '94 Symposium, pp. 24-35, Aug 1994.
- [145] Movie Trace Files. <http://www-tnk.ee.tu-berlin.de/research/trace/ltvt.html>
- [146] J. Farber, "Traffic Modelling for Fast Action Network Games," Multimedia Tools and Applications, vol. 23, no. 1, pp. 31-46, 2004.
- [147] IEEE Standard for Information Technology, "Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks - Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment: Medium Access Control (MAC) Quality of Service Enhancements", P802.11e/D13.0, 2005.
- [148] G. Marfia, C. E. Palazzi, G. Pau, M. Gerla, M. Y. Sanadidi, M. Roccetti, "TCP Libra: Exploring RTT-Fairness for TCP", UCLA CSD Technical Report #TR050037, 2005.
- [149] Consumer Electronics - IT Facts. <http://www.itfacts.biz/index.php?id=P1861>

- [150] InternetFrog.com. <http://www.internetfrog.com/myipc/speedtest/>
- [151] Bandwidth Speed Test. <http://www.bandwidthplace.com/speedtest/>
- [152] Broadband Speed Tests. <http://www.dslreports.com/stest>
- [153] Verizon Online DSL. <http://www.verizon.com/dsl/>
- [154] Bellsouth Fastaccess DSL <https://www.fastaccess.com/content/splash/>
- [155] AT&T Worldnet DSL Service. <http://www.att.net/dsl/>
- [156] H. Jiang, C. Dovrolis, "Why is the Internet traffic bursty in short (sub-RTT) time scales?", in Proc. of ACM SIGMETRICS 2005, Banff, AL, Canada, 2005.
- [157] IEEE, "Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," Specifications, ISO/IEC 8802-11:1999(E), 1999.
- [158] A. Balk, M. Gerla, M. Sanadidi, D. Maggiorini, "Adaptive Video Streaming: Pre-encoded MPEG-4 with Bandwidth Scaling", Computer Networks: The International Journal of Computer and Telecommunications Networking, vol. 44, no. 4, pp. 415-439, 2004.
- [159] J. Padhye, V. Firoiu, D. Towsley, J. Kurose, "Modeling TCP Throughput: A Simple Model and its Empirical Validation", in Proc. of the ACM SIGCOMM '98 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, pp. 303-314, Vancouver, B.C., Canada, 1998.
- [160] C. Jin, D. Wei, S. H. Low, G. Buhmaster, J. Bunn, D. H. Choe, R. L. A. Cottrell, J. C. Doyle, W. C. Feng, O. Martin, H. Newman, F. Paganini, S. Ravot, and S. Singh, "Fast TCP: From Background Theory to Experiments," IEEE Network, vol. 19, no. 1, pp. 4-11, 2005.
- [161] D.-Y. Chen, S. Garg, M. Kappes, K. S. Trivedi, "Supporting VBR VoIP Traffic with IEEE 802.11 WLAN in PCF Mode", in Proc. of OPNETWork 2002, Washington D.C., 2002.
- [162] D. Chen, D. Gu, J. Zhang, "Supporting Real-time Traffic with QoS in IEEE 802.11e Based Home Networks", Mitsubishi Electric Research Laboratory, TR-2004-006, 2004.
- [163] C. E. Palazzi, S. Ferretti, S. Cacciaguerra, M. Roccetti, "A RIO-like Technique for Interactivity Loss Avoidance in Fast-Paced Multiplayer Online Games", ACM Journal of Computers in Entertainment, vol.3, no.2, Apr 2005.
- [164] S. Ferretti, C. E. Palazzi, M. Roccetti, G. Pau, M. Gerla, "FILA, a Holistic Approach to Massive Online Gaming: Algorithm Comparison and Performance Analysis", in Proc. of the 3rd ACM International Conference in Computer Game Design and Technology (GDTW 2005), Liverpool, UK, pp. 68-76, Nov 2005.
- [165] S. Ferretti, C. E. Palazzi, M. Roccetti, G. Pau, M. Gerla, "FILA, a Holistic Approach to Massive Multiplayer Online Gaming: Algorithm Comparison and Performance Analysis", *to appear* in the ACM Journal of Computers in Entertainment.
- [166] C. E. Palazzi, G. Pau, M. Roccetti, M. Gerla, "In-Home Online Entertainment: Analyzing the Impact of the Wireless MAC-Transport Protocols Interference", in Proc. of IEEE International Conference on Wireless Networks, Communications and Mobile Computing (WIRELESSCOM 2005), Maui, HI, USA, Jun 2005.
- [167] C. Gauthier Dickey, D. Zappala, V. Lo, J. Marr, "Low Latency and Cheat-Proof Event Ordering for Peer-to-Peer Games", in Proc. of ACM International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV), Kinsale, County Cork, Ireland, pp. 134-139, Jun 2004.

- [168] N. E. Baughman, B. N. Levine, "Cheat-proof Payout for Centralized and Distributed Online Games", in Proc. of IEEE Infocom, Anchorage, AK, USA, pp. 104-113, Apr 2001.
- [169] Gamasutra. http://www.gamasutra.com/features/20000724/pritchard_01.htm
- [170] B. D. Chen, M. Maheswaran, "A Fair Synchronization Protocol with Cheat Proofing for Decentralized Online Multiplayer Games", in 3rd IEEE Symposium on Network Computing and Applications (NCA), Cambridge, MA, USA, pp. 372-375, Aug 2004
- [171] W. Bernier, "Latency Compensating Methods in Client/Server In-game Protocol Design and Optimization", in Game Developer Conference (GDC), Mar 2001. <http://www.resourcecode.de/stuff/clientsideprediction.pdf>
- [172] M. Bottigliengo, C. Casetti, C. Chiasserini, M. Meo "Short-Term Fairness for TCP Flows in 802.11b WLANs", in Proc. of IEEE INFOCOM 2004, Hong Kong, China, Mar 2004.
- [173] S. Garg, M. Kappes, "An Experimental Study of Throughput for UDP and VoIP Traffic in IEEE 802.11b Networks", IEEE Wireless Communications and Networking Conference (WCNC 2003), New Orleans, LA, USA, pp. 1748-1753, Mar 2003.
- [174] A. L. Wijesinha, Y. Song, M. Krishnan, V. Mathur, J. Ahn, V. Shyamasundar, "Throughput Measurement for UDP Traffic in an IEEE 802.11g WLAN", in Proc. of 6th International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing and First ACIS International Workshop on Self-Assembling Wireless Networks (SNPD/SAWN'05), Towson, MD, USA, pp. 220-225, May 2005.
- [175] D. P. Hole, F. A. Tobagi, "Capacity of an IEEE 802.11b Wireless LAN supporting VoIP", in Proc. IEEE International Conference on Communications (ICC 2004), Paris, France, pp. 196-201, Jun 2004.
- [176] G. Bianchi, "IEEE 802.11-Saturation Throughput Analysis" IEEE Communications Letters, vol. 2, no. 12, pp. 318-320, Dec 1998.
- [177] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function", IEEE Journal on Selected Areas in Communications, vol. 18, no. 3, pp.535-547, Mar 2000.
- [178] N. Shachtman, "New Army Soldiers: Game Gamers", Wired News, Oct 2001. <http://www.wired.com/news/conflict/0,2100,47931,00.html>
- [179] Rakion. <http://www.rakion.com/>