

WPI-CS-TR-21-06

October 2021

A Survey and Taxonomy of Latency Compensation Techniques for Network
Computer Games

by

Shengmei Liu
Xiaokun Xu and Mark Claypool

For latest version and interactive taxonomy, see:
<https://web.cs.wpi.edu/~claypool/papers/lag-taxonomy/>

Computer Science
Technical Report
Series



WORCESTER POLYTECHNIC INSTITUTE

Computer Science Department
100 Institute Road, Worcester, Massachusetts 01609-2280

A Survey and Taxonomy of Latency Compensation Techniques for Network Computer Games

SHENGMEI LIU, XIAOKUN XU, and MARK CLAYPOOL, Worcester Polytechnic Institute, USA

Computer games, one of the most popular form of entertainment in the world, are increasingly online multi-player, connecting geographically dispersed players in the same virtual world over a network. Network latency between players and the server can decrease responsiveness and increase inconsistency across players, degrading player performance and quality of experience. Latency compensation techniques are software-based solutions that seek to ameliorate the negative effects of network latency by manipulating player input and/or game states in response to network delays. We search, find, and survey over 80 papers on latency compensation, organizing their latency compensation techniques into a novel taxonomy. Our hierarchical taxonomy has eleven base technique types, organized into four main groups. Illustrative examples of each technique are provided as well as demonstrated use of the techniques in commercial games.

CCS Concepts: • **Applied computing** → **Computer games**.

Additional Key Words and Phrases: video games, lag, quality of experience, user study, delay

ACM Reference Format:

Shengmei Liu, Xiaokun Xu, and Mark Claypool. 2022. A Survey and Taxonomy of Latency Compensation Techniques for Network Computer Games. *ACM Comput. Surv.* 1, 1, Article 1 (January 2022), 35 pages. <https://doi.org/10.1145/3519023>

1 Introduction

The computer games market was worth over \$150 USD globally in 2019 and is expected to grow about 13% per year from 2020 to 2027 [57]. There are about 2.7 billion computer gamers world-wide [67] who are increasingly playing games online. Online games are estimated to reach \$196 billion USD in revenue by 2022 and are one of the fastest growing industries in the world [72].

While many factors can affect a player's gameplay quality of experience, multiplayer network games – where physically separated players simultaneously interact in a shared, virtual game world – must overcome the challenges posed by computer networks: limited bitrates, lost game data from dropped packets, and latency from sending game information from one computer to another. While increased capacities can overcome bitrate limitations and reduce packet loss, latency remains a challenge for network games and can never be eliminated entirely. Latency determines not only how players experience gameplay but also how network games must be designed in order to mitigate the effects of latency and meet player expectations.

Authors' address: Shengmei Liu, sliu7@wpi.edu; Xiaokun Xu, xxu11@wpi.edu; Mark Claypool, claypool@cs.wpi.edu, Worcester Polytechnic Institute, Worcester, MA, USA, 01609.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

Unchecked, latencies can degrade both player performance (e.g., lower game score) and quality of experience (e.g., less fun).

For multiplayer network games, the canonical game state is typically controlled by an authoritative server, and each player is on a separate client computer with a copy of the game state. The clients capture and send player actions to the server which updates the game world appropriately and sends the new game state back to each client. The time it takes from when the player provides input until the new game state is rendered is *latency* in this paper (colloquially called “lag” by gamers). Thus, a primary impact of latency is degraded *responsiveness* for the player in that there is an observable delay between player input and game state update. In addition, player input and game update messages take time to propagate and process before all client and server states are the same. Thus, another impact of latency is reduced *consistency* between the game states at the clients and the game state at the server.

Latency compensation techniques are software algorithms that run on the game client or game server (or both) that try to mitigate the impact of latency on the players. Such techniques might improve responsiveness or increase consistency (or both), but often sacrifice one for the other (i.e., increased consistency at the cost of reduced responsiveness and vice versa). The effectiveness of the techniques depends upon many factors including, but not limited to, the game applied to, the network conditions between clients and servers, and the familiarity and skill of the players with the game.

Understanding what latency compensation techniques are effective and for what games and which network and game conditions can be helpful for game developers and game system developers that want to improve the multi-player network game experience, and for researchers that seek to improve upon existing techniques. Such understanding can be gained through a *survey* of latency compensation techniques that collates information to provide an overview of existing techniques. Or, even better, through a *taxonomy* that categorizes latency compensation techniques to illustrate the relationships different techniques have in common and organize the information so as to make accumulated survey information accessible based on a user’s need.

While there has been one survey of latency compensation techniques by Jiang et al. [71], it provides for only two types of latency compensation techniques (synchronization and optimistic) and is 15 years old so may not represent techniques developed since. An early, often cited de facto survey of latency compensation techniques by Bernier [10] describes a few techniques in detail, but is 20 years old and does not encompass the full set of techniques developed today nor describe their use in modern games. A more recent survey by Briscoe et al. [14] surveys techniques to reduce network latency but includes only techniques at the systems level and without specific attention to computer games. We are not aware of any other surveys or taxonomies for latency compensation techniques for network games.

This work provides a survey in the form of a thorough literature review of peer-reviewed publications on latency compensation techniques for computer games and similar interactive media. We began with a core set of known (to us) papers on latency compensation techniques, then expanded our pool via following forward and backward (in time) citations to these papers and searching Google Scholar. This yielded over 80 peer-reviewed publications that dealt with latency compensation for games. We used these 80+ papers as the foundation for a novel taxonomy of latency compensation techniques, divided into 4 main groups, each further broken down into 11 final categories.

Our survey further identifies the games, game genres and methods of evaluation for each paper, showing over half of the papers do not include evaluations with users and about 20% have no evaluation at all. As shown by others [106, 120, 134, 149], latency compensation algorithms must be tested and then evaluated to determine their effect. Thus, our work also provides an indication as to where additional work might be done to assess the efficacy of techniques over a broader range of games and game conditions.

It is believed that many latency compensation techniques are implemented into modern multiplayer network games. However, as is typical of much commercial software, computer games are usually closed-source and implementations are not specified. However, we have located about a dozen blogs from commercial game developers where they indicate latency compensation approaches implemented for their games. We have categorized these using our taxonomy and provided links as specific evidence of latency compensation techniques used in practice.

The main contributions of this work include:

- Examples illustrating the main effects of latency on multiplayer network games: *responsiveness* and *consistency*.
- A collection of 82 peer-reviewed research papers that deal with latency compensation, providing information on type of technique, evaluation method, game and game genre evaluated and latency range studied,
- A visual depiction of a taxonomy for latency compensation techniques, showing a hierarchical representation of techniques grouped by similarity.
- Definitions for each technique (node) in the taxonomy.
- A visual example for each “leaf” in the taxonomy, illustrating the application of a taxonomy to a computer game.
- Identification of 20 presentations and blogs by commercial game developers detailing the latency compensation techniques applied to their games.

Section 2 provides background information on latency, interactivity and games; Section 3 describes our methodology to survey peer-reviewed research papers on latency compensation techniques; Section 4 presents our taxonomy; Section 5 lists blogs that provide evidence of latency compensation in commercial games; Section 6 discusses our insights provided by the survey and taxonomy; and Section 7 summarizes our paper and presents possible future work.

2 Latency, Interactivity and Games

This section provides background information on latency in computer game systems (Section 2.1), the effects of latency on computer games (Section 2.2), and client-server game architectures (Section 2.3).

2.1 Latency in Computer Game Systems

In general, latency in a computer game affects how long it takes from when a player acts until the player sees/hears the results. And there is *always* some latency between player actions and game output.

Latencies around 20-30 milliseconds can be noticed for interactive music [101], and even lower at around 10 milliseconds [68] and perhaps even 2 milliseconds [102] for dragging-related tasks. Latency around 50 ms is known to affect performance in mouse-based pointing tasks [68, 95]. Latencies under 100 milliseconds

have been demonstrated to affect game-related tasks, too, such as moving target selection and steering [52, 94, 107].

For a non-network game, the latency is from the player’s local computer and includes delays from processing by the input devices, operating system, game and game engine, graphics cards and display devices. Such latency is sometimes called “click to photon” latency by gamers, referring to the time from a mouse click until the monitor displays the result. Local latencies for games can range from high-performance game systems with latencies around 25 milliseconds [93] to console and TV combinations that have latencies of over 200 milliseconds [66, 110].

For network games, the local system latency still impacts the player, but there is also *network latency* – the round-trip time between the game client and the game server, including all network processing on the end-hosts (the network interface cards and OS stack) and all the intermediate devices/hops between. The primary focus of latency compensation techniques is to mitigate these network latencies. Network latencies can vary by several orders of magnitude, from milliseconds for a Local Area Network (LAN), to 10s of milliseconds for an intra-continental network, to 100s of milliseconds for a inter-continental network. Wireless networks, such as WiFi, mobile and satellite, can even experience seconds of latency under poor network conditions.

For a network game, the player is vulnerable to *both* network latency and local latency. However, most latency compensation techniques are designed to overcome network latency and not local latency.

2.2 The Effects of Latency on Network Games

Latency can make a network computer game *less responsive*. Consider the example in Figure 1a showing an event timeline for a network game with a client and a server, with time progressing top to bottom. The left side shows the display on the client over this time. In this example, the player provides input and the client encapsulates that input into a message that is sent to the server. The server receives the message, processes the input, updates the game world, and sends the world update back to the client. The client, upon getting the update, renders the new world for the player on the display. The time from the player input until the resulting change to the display is the response time. The larger the network latency, the greater the response time and the less responsive the computer game.

Latency can also make a network computer game *less consistent*. Consider the example in Figure 1b showing a multi-player networked game with two clients with different latencies connected to a server. The authoritative server has the true location of an avatar depicted by the green circle. The avatar is moving up and two the right along the path. The server periodically sends position updates to each client. Ideally, in most game condition, such as a race where avatars are running together around a virtual track, the view of the game world on both Client A and Client B would be the same (consistent), but because of latency they are not. Since Client A is 25 milliseconds away from the Server, the displayed position is close to that of the server. However, since Client B is more distant from the Server at 100 milliseconds, the displayed position is even farther away from the actual position. The client game worlds are inconsistent with each other and the server.

This impact of latency can be visualized as in Figure 2, where consistency and responsiveness are depicted by orthogonal axes. In this space, the best position for a network game – the one with the highest responsiveness and most consistency – is at the origin at the bottom left. The further away from the origin,

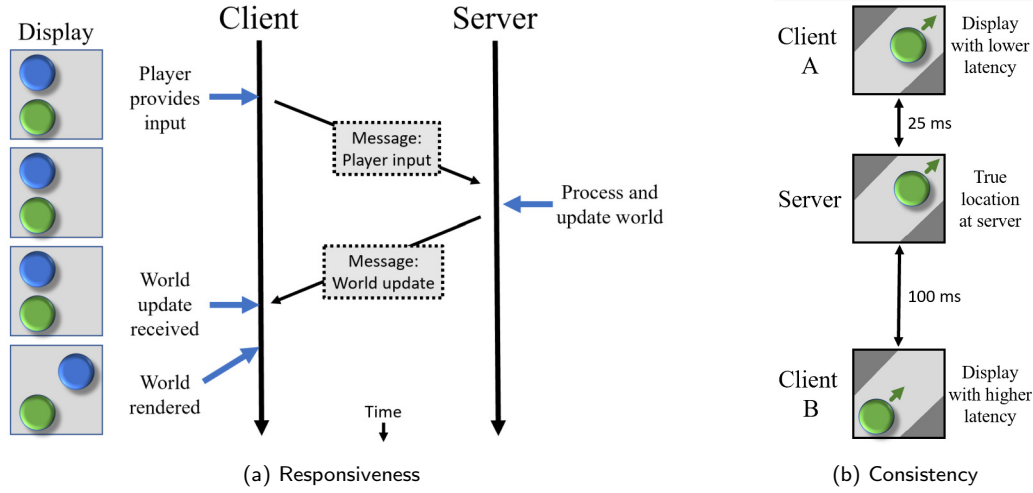


Fig. 1. The effects of latency on network games.




Fig. 2. Tradeoff between consistency and responsiveness.

the worse (i.e., the less responsive and/or more inconsistent). Client A from the example in Figure 1a might appear at point *A* somewhat near the origin, with Client B at point *B* being further away. Latency compensation techniques seek to improve the responsiveness or consistency (or both) in the presence of latency, moving a client closer to the origin. For example, one technique may improve the consistency of Client B to position *B'*, while another may improve the responsiveness of Client A to position *A'*, at the cost of a decrease in consistency.

Because there is always at least some latency with a network computer game, being at the origin is impossible. In fact, there is often a tradeoff between consistency and responsiveness. Some latency compensation techniques designed to improve responsiveness in the presence of latency can decrease consistency, while other techniques that improve consistency can make a game less responsive.

There has been numerous studies on the effects of network latency and games, assessing the effects of latency on both player performance and quality of experience (QoE).

For first-person shooter (FPS) games, Armitage et al. [7] estimate a player tolerance threshold for latency for Quake 3 is about 150-180 ms. Dick et al. [36] show via a survey that players generally think about 120 ms is the maximum tolerable latency for a network game, regardless of game genre, but their user study shows players find 150 ms acceptable for the games Counter-strike and Unreal Tournament 2003 (UT2003). Quax et al. [109] find for UT2003 players, latency and latency jitter as low as 100 ms can degrade player performance and quality of experience. Amin et al. [4] show player experience determines the sensitivity to latency for Call of Duty, with competitive gamers more adept at compensating for impaired conditions. Liu et al. [92] for Counter Strike: Global Offensive find decreasing latencies from 125 milliseconds to 25 milliseconds improves player accuracy by about 2%, score by about 15% and QoE by about 20%.

For other genres, Pantel and Wolf [105] show latencies of about 100 ms can affect car racing games. Fritsch et al. [53] find players of the role-playing game Everquest 2 can tolerate hundreds of milliseconds of network latency. Hohlfeld et al. [63] find players of the casual game Minecraft are insensitive to network latencies of up to 1 second. Sheldon et al. [123] show some aspects of play in the real-time strategy game Warcraft 3 are not affected by even a second of latency. Howard et al. [64] indicate that cooperation in the role-playing game Mass Effect 3 can be affected by latency of a teammate due to cascading effects on the game outcome.

Many of these studies use specific games, but are intended to generalize to games from the same genres (e.g., first-person shooter games) [4, 7, 9, 22, 25]. However, studies using commercial games include a game engine’s built-in latency compensation techniques, often confounding the results.

Generally, the effects of latency depend upon the game type, the player actions in the game and the characteristics of those actions [26, 114]. For example, turn-based games are more forgiving of latency than are games with tight interactivity. Based on some of this work, Claypool and Claypool [26] suggest a game action’s sensitivity to latency can be classified by precision and deadline – higher precision and tighter deadline mean more sensitivity to latency. For game combat, shooting with a high-precision weapon is affected by latency more than shooting with a weapon with a wide area of effect. For game navigation, driving a car around a race track is affected by latency more than sending an avatar across a large world to explore. Additional research suggests the effects of latency also depend upon a player’s skill, with higher skill players more affected by latency than lower skill players [92].

2.3 Client-Server Game Architectures

While underlying network game architectures can be peer-to-peer [126], most multiplayer network games use a client-server architecture. This architecture is popular for network games since firewall rules can make it difficult for clients to connect to each other. Moreover, having a single server can help a game scale with number of players. Last but not least, architectures with a trusted authoritative server (e.g., managed by the game publisher) can help thwart player cheating since the server keeps the official game state and makes all final decisions about the outcome of player actions.

Typically, the server is at a well-known IP address and port and is public so as to be reachable by all clients playing together in one game. In some cases, a server browser setup allows game servers to register their individual IP addresses and ports with a central server, allowing clients to connect to the central server and browse available game servers, the individual game servers being differentiated based on configuration parameters (e.g., a certain game mode, map or network latency). Once an individual game

server is chosen, the central server provides the client with the game server IP address and port whereupon the client connects to the game server to play the game.

Some network game architectures have one client act as the host to which the other clients connect. While at the network level, these have peer-to-peer connections (i.e., one client communicates directly with another client), they can still be viewed as have a client-server architecture in that the host acts as a server, albeit one that also acts as a client for that player.

3 Methodology

Our methodology is designed to survey peer-reviewed research papers on latency compensation techniques applied to network games and similar interactive applications.

A comprehensive literature review survey examines all (or nearly all) scholarly sources related to a specific topic or research question (in our case, latency compensation). Unfortunately, given the general nature of representative keywords, doing an exhaustive search via the Web is prohibitive. For example, as of November 2020, a search in Google Scholar with the keywords “Latency Compensation” had 144,000 results, “Latency Compensation and Interactive Media” had 24,800 results and “Latency Compensation and Games” had 11,600 results.

Instead of a top-down search approach, we opted for a bottom-up approach. We started with about 20 initial research papers that dealt with latency compensation. Careful review of these papers provided the initial version of the taxonomy. From this list of papers, we expanded our reference list based on earlier works the original papers cited as well as forward (in time) references that cited these papers. Our main selection criteria were papers that dealt with latency compensation in some significant fashion. For example, a paper that proposed, discussed or evaluated a latency compensation technique was included, while a paper that merely mentioned latency compensation in the related work section was not included.

This yielded a pool of about 50 papers.

From there, we did a search in Google Scholar with the keywords “Latency Compensation and Games”, sorted them by relevancy, and went through the first 200 entries, adding those that involved latency compensation and we did not yet have in our pool. We repeated these same steps for the keywords “Delay Compensation and Games”.

For each paper in the pool, we categorized the work based on the initial taxonomy and recorded relevant details on the paper. We also noted equivalency terms used in the paper in order to map them to the terms used in our taxonomy. Papers were discussed by us to confirm details and placement in the taxonomy. Whenever appropriate (i.e., a paper’s technique did not fit into our taxonomy), we would rework categories in the taxonomy until the technique could be placed. We also examined the list of works cited and added any missing previous references to our pool. We continued this process until the pool was exhausted (i.e., we had reviewed and classified all latency compensation papers known to us).

For inclusion in our list, we only considered peer-reviewed publications (i.e., not white papers, technical reports, theses nor patents) and made no determination of the “quality” of the peer-reviewed forum nor the quality of the paper itself. When in doubt (i.e., we were not sure if the paper was peer-reviewed or invited), we included the paper.

For topics, we only consider multiplayer interactive systems, so this excludes the related area of tele-operation, which can also suffer with latency. The interested reader is encouraged to refer to a survey by Farajiparvar et al. [44] for more information.

We also exclude literature that only focuses on local system latency (i.e., not network latency), such as those that predict mouse movements [20] or hardware upgrades that might reduce local latency, such as monitors with low refresh rates [130], and computer mice with high polling rates. We also do not consider latency reduction in low-level controllers, such as latency compensation for power oscillation damping controllers [21].

Network system techniques for reducing latency are also excluded, including but not limited to choice of transport protocol, upgrades to link capacities, sizing of router buffers and path routing. The interested reader is encouraged to refer to a survey by Brisco et al. [14].

Similarly, we also exclude techniques that reduce latency through system reconfiguration – e.g., moving a server closer to the player. While server placement can have a marked impact on client-server latency, such techniques are not done in-game and must be done *a priori*. Moreover, moving servers often cannot be done for multi-person games where players are geographically separated [55] or for games deployed without the resources needed for a distributed architecture.

Instead, the techniques surveyed assume there is a base level of latency present in the game system (i.e., local latency), and there is an unavoidable network latency between the clients and the server. The latency compensation techniques seek to mitigate the network latency by end-host computations on the clients or the server (or both), adjusting anything between player input and game actions to game state and rendered output.

4 Survey and Taxonomy

Our methodology yielded 82 peer-reviewed publications, surveyed and taxonomized in this section. Table 1 provides the list of papers.

Year is when the paper was published.

Compensation Technique maps the latency compensation technique(s) in the paper to our taxonomy. Note, many papers include more than one technique.

Study Type is one of: “User Study” involving human subjects, “Experiments” using games or simulated games but no active human players, or “Case Study” with a single instance described as a proof of concept. If the paper has no meaningful evaluation the study type is listed as “None”.

Game provides the name(s) of the game(s) used in evaluation (if done). Game names are provided for commercial games or publicly available games – custom-built games are so indicated as “Custom”. Studies evaluating a task rather than a game – e.g., drawing [6] or writing [97] – are listed as “Task”.

Genre is the general class to which the evaluated games belong: “FPS” (first-person shooter), “RPG” (role playing game), “RTS” (real-time strategy), “Casual” (also covering turn-based games), “Fighting”, “Racing”, “Sports” and “Arcade”. The Arcade genre is a bit of a catch-all covering games with limited game actions (and gameplay depth) and simple graphics. For cases where the evaluation is an experiment with generic interaction, the genre is listed as “Any”.

Range indicates the latency values evaluated in papers where evaluation is done and the latencies specified.

Users has the number of users in the evaluation, where applicable.

Table 1. Peer-reviewed publications with research on latency compensation for network games.

Year	Paper	Compensation Technique	Study Type	Game	Genre	Range	Users
1985	[70]	Time Warp	None	N/A	N/A	N/A	N/A
1992	[129]	Latency Concealment	User Study	Custom	N/A	0 - 380 ms	12
1994	[100]	Extrapolation	None	N/A	N/A	N/A	N/A
1998	[122]	Interpolation	Case Study	Custom	Arcade	0 - 2000 ms	2
1999	[17]	Extrapolation, Latency Exposure	Experiment	N/A	N/A	N/A	N/A
1999	[139]	Attribute Scaling, Extrapolation	None	N/A	N/A	N/A	N/A
1999	[38]	Latency Exposure					
1999	[38]	Extrapolation, Incoming Delay	User Study	MiMaze	Arcade	0 - 100 ms	25
2000	[98]	Extrapolation, Time Warp	None	N/A	N/A	N/A	N/A
2000	[51]	Latency Exposure	None	N/A	N/A	N/A	N/A
2000	[142]	Self-Prediction	User Study	Custom	Arcade	120 - 300 ms	8
2000	[97]	Extrapolation, Incoming Delay, Time Warp	Case Study	Task	N/A	0 - 300 ms	N/A
2001	[121]	Extrapolation	None	N/A	N/A	N/A	N/A
2001	[145]	Extrapolation	Experiment	Custom	Any	unspecified	N/A
2002	[106]	Extrapolation	Experiment	Custom	Arcade, Racing	100 - 200 ms	N/A
2002	[127]	Extrapolation	None	N/A	Sports	N/A	N/A
2002	[91]	Incoming Delay	Experiment	Custom	FPS	0 - 300 ms	N/A
2003	[30]	Extrapolation	None	N/A	N/A	N/A	N/A
2003	[59]	Extrapolation	User Study	Task	N/A	0 - 480 ms	8, 18
2003	[41]	Extrapolation	Experiment	Custom	Arcade	unspecified	N/A
2004	[1]	Extrapolation	Case Study	BZ Flag	FPS	100 - 800 ms	2, 4
2004	[99]	Incoming Delay, Time Warp	Experiment	Custom	Arcade	40 - 2000 ms	N/A
2004	[128]	Interpolation	None	N/A	N/A	N/A	N/A
2004	[58]	Latency Exposure	User Study	Custom	Arcade	0 - 1400 ms	40
2004	[147]	Extrapolation	Experiment	Task	N/A	unspecified	N/A
2005	[2]	Extrapolation	Experiment	BZ Flag	FPS	200 - 800 ms	4
2005	[74]	Latency Concealment, Self-Prediction	None	N/A	N/A	N/A	N/A
2005	[124]	Extrapolation, Latency Concealment	Case Study	Custom	Arcade	0 - 200 ms	N/A
2005	[146]	Self-Prediction					
2005	[146]	Outgoing Delay	Experiment	Quake 2	FPS	0 - 400 ms	4
2006	[15]	Extrapolation, Incoming Delay, Outgoing Delay, Self-Prediction	Experiment	Custom	Any	0 - 500 ms	N/A
2006	[71]	Extrapolation, Time Warp	None	N/A	N/A	N/A	N/A
2006	[16]	Latency Concealment, Self-Prediction	Experiment	Custom	Casual	N/A	N/A
2006	[141]	Extrapolation, Latency Exposure	User Study	Custom	Sports	50 - 250 ms	24
2006	[61]	Outgoing Delay					
2006	[61]	Extrapolation	Experiment	Task	N/A	N/A	4
2006	[148]	Extrapolation, Incoming Delay	User Study	Custom	Arcade	100 - 800 ms	2
2006	[87]	Extrapolation	User Study	Custom	Arcade	0 - 3000 ms	30
2006	[104]	Extrapolation	Experiment	BZ Flag	FPS	0 - 1600 ms	N/A
2007	[23]	Incoming Delay, Self-Prediction	User Study	Task	N/A	0 - 900 ms	18
2007	[79]	Incoming Delay	User Study	Custom	FPS	25 - 200 ms	10
2007	[88]	Extrapolation	User Study	Custom	Arcade	1000 - 3000 ms	37
2007	[137]	Self-Prediction	Experiment	Task	N/A	0 - 50 ms	N/A
2008	[134]	Incoming Delay	User Study	Custom	Arcade	0 - 500 ms	24
2008	[113]	Extrapolation	User Study	Custom	Arcade	N/A	3
2008	[96]	Control Assistance	User Study	Custom	Arcade	N/A	13
2008	[89]	Extrapolation	User Study	Custom	Arcade	0 - 3000 ms	37
2008	[103]	Incoming Delay	Experiment	Custom	Any	2 - 40 ms	N/A
2009	[149]	Incoming Delay	Experiment	Street Fighter 2	Fighting	0-400 ms	2
2009	[65]	Extrapolation, Incoming Delay	User Study	Custom	Racing	unknown	unknown
2010	[120]	Extrapolation, Incoming Delay, Outgoing Delay, Self-Prediction, Time Warp	Experiment	Custom	Arcade	100 - 746 ms	N/A
2010	[73]	Outgoing Delay	Experiment	Quake 3 Arena	FPS	0 - 150 ms	N/A
2011	[8]	Control Assistance	User Study	Custom	Arcade	N/A	24
2012	[75]	Extrapolation	Experiment	Custom	Any	unspecified	N/A
2012	[62]	Extrapolation, Incoming Delay	User Study	Custom	Arcade	0-150	30
2012	[78]	Extrapolation, Incoming Delay	User Study	Custom	Arcade	80, 130 ms	20
2013	[119]	Extrapolation, Incoming Delay	None	N/A	N/A	N/A	N/A
2013	[144]	Interpolation					
2013	[144]	Extrapolation	User Study	Quake 3, World of Warcraft	FPS, RPG	N/A	16, 200
2013	[125]	Attribute Scaling	None	Custom	Sports	100 - 200 ms	N/A
2013	[143]	Attribute Scaling, Incoming Delay	User Study	Custom	Fighting	10 - 40 ms	20
2014	[117]	Interpolation	User Study	Custom	Arcade	125 - 500 ms	18
2014	[118]	Extrapolation, Incoming Delay	User Study	Custom	Arcade	50 - 200 ms	26
2014	[140]	Control Assistance	User Study	Custom	FPS	N/A	N/A
2015	[66]	Control Assistance	User Study	Custom	Arcade	10 - 160 ms	18
2015	[83]	Speculative Execution	User Study	Doom 3, Fable 3	FPS, RPG	0 - 400 ms	41
2015	[86]	Extrapolation, Interpolation, Time Warp	User Study	CS:GO	FPS	0 - 150 ms	4
2016	[54]	Extrapolation	Experiment	Custom	Arcade	N/A	N/A
2016	[138]	Self-Prediction	Experiment	N/A	N/A	50 - 100 ms	N/A
2016	[43]	Latency Concealment	None	N/A	N/A	N/A	N/A
2016	[69]	Extrapolation	Experiment	Custom	Racing	N/A	15
2017	[84]	Time Warp	User Study	Custom	FPS	0 - 250 ms	4
2017	[80]	Self-Prediction	User Study	Task	N/A	33 - 99 ms	16
2018	[90]	Outgoing Delay	User Study	Assault Cube	FPS	20 - 80 ms	10
2018	[114]	Attribute Scaling	User Study	Need for Speed, Somi	Arcade, Racing,	0 - 400 ms	25, 27
2018	[85]	Table Tennis					
2018	[85]	Time Warp	User Study	Custom	FPS	50 - 250 ms	12
2018	[32]	Extrapolation	Experiment	Custom	Arcade	0 - 300 ms	N/A
2018	[24]	Extrapolation	Experiment	Custom	Racing	N/A	N/A
2018	[6]	Self-Prediction	User Study	Task	N/A	0 - 66 ms	10
2019	[82]	Attribute Scaling	User Study	Flappy Bird	Arcade	0 - 200 ms	12
2019	[135]	Time Warp	User Study	Custom	Arcade	0 - 800 ms	30
2020	[40]	Extrapolation	Experiment	World of Warcraft	RPG	N/A	N/A
2020	[116]	Attribute Scaling	User Study	Flappy Bird	Arcade	10 - 400 ms	18
2020	[115]	Attribute Scaling, Control Assistance	User Study	Custom	Arcade	0 - 400 ms	9
2020	[76]	Latency Concealment	User Study	Custom	FPS	25 - 105 ms	9
2021	[12]	Latency Concealment	Case Study	Custom	Arcade	unspecified	N/A

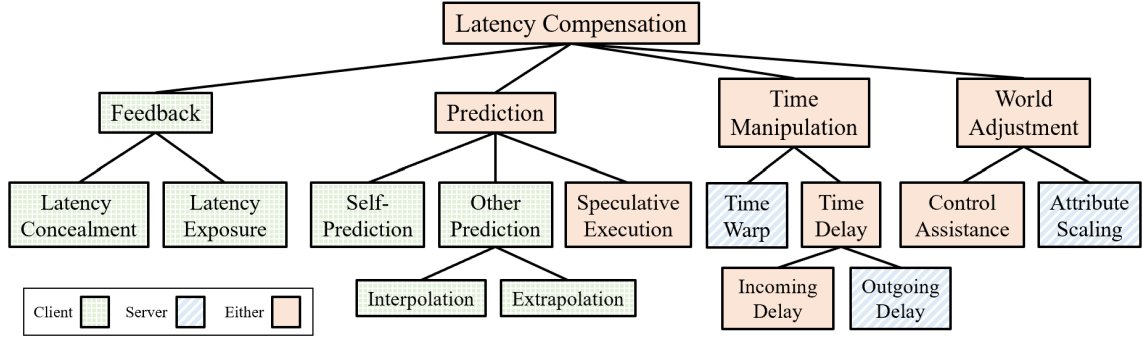


Fig. 3. A taxonomy of latency compensation techniques for network games.

Figure 3 depicts our taxonomy of latency compensation techniques for network games. The techniques are arranged in groups based on shared characteristics. The proposed taxonomy is hierarchical, read from the top down, where each vertical level provides increased differentiation and refinement of the indicated technique – i.e., more general on the top to more specific on the bottom.

At the top is *Latency Compensation*, grouping all latency compensation techniques.

Beneath *Latency Compensation* is the division of 4 classes of latency compensation: *Feedback*, *Prediction*, *Time Manipulation* and *World Adjustment*.

Each of these classes is further differentiated by families of techniques: *Latency Concealment* and *Latency Exposure* for *Feedback*, *Self-Prediction*, *Other-Prediction*, and *Speculative Execution* for *Prediction*, *Time Warp* and *Time Delay* for *Time Manipulation*, and *Control Assistance*, and *Attribute Scaling* for *World Adjustment*.

Some families of techniques are further refined: *Interpolation* and *Extrapolation* for *Other-prediction* and *Incoming Delay* and *Outgoing Delay* for *Time Delay*.

Although not shown in the figure, actual implementations of latency compensation techniques can also be differentiated based on algorithm parameters or code specifics.

Figure 3 also indicates by color and shading where the primary functionality for the compensation technique resides (see the legend in the bottom left corner): *Client* for techniques that are processed and displayed primarily on a player’s client computer, *Server* for techniques that are primarily handled on the authoritative game server, and *Either* for techniques that can reside on either the client or the server or both.

Note, individual network games can, and often do, use more than one technique.

The rest of this section proceeds to define and describe the nodes in the taxonomy, providing an illustrative example for each “leaf” node in the picture. The examples are not necessarily the only latency compensation technique in each leaf node, but are meant to provide clarity by an example.

4.1 Feedback

Feedback provides audible or visual information to the player based on latency, without actually changing the state of the game world.

4.1.1 Latency Concealment visually masks latency from the client to the server so as to minimize the perception of unresponsiveness. For example, a game client might show the discharge and recoil of a weapon when the player pulls the trigger, even though the outcome of the shot (e.g., a hit or a miss) has not been registered and recorded by the server. As another example, a vehicle ordered to move by a player may immediately power up its engines, kicking up dust but not actually change positions until the server has verified movement is allowed.

Figure 4 depicts an example of latency concealment. On the left, at time t_0 the player is ready to act to move their avatar, the green circle. At time t_1 when the player provides input at the game client, the game responds by providing a visual effect (if this was a car, the engines might rev, kicking up dust), but the avatar does not actually change positions pending approval from the authoritative server. At time t_2 , the avatar actually moves since the server approval has arrived. The visual effect in this example conceals the latency from the client to the server by making the move command seem immediately responsive to the player, even though the actual position does not change until at least as long as the client-server latency.




Fig. 4. An example of latency concealment.

Latency concealment includes image warping techniques which adjust the rendered image, ranging from simple camera view re-projection [129] to temporally updating the animated content [12, 43, 74, 76]. These concealment techniques are often used in virtual reality headsets (called *late warp*) to reduce latency for head motion to video output to help with simulator sickness, and other aspects.

Burgess, Hanna, Shelly and Katchabaw [16, 74, 124] describe the use of a software design pattern that uses “padding” to hide latency. They make specific mention of using an animation played by the client in immediate response to a player’s action while waiting for the latency from the server confirmation.

4.1.2 Latency Exposure gives a visual indicator of the magnitude of the latency from the client to the server. For example, the client may decorate the corner of a display with a numeric value that directly reports the client’s round-trip time to the server (often called the “ping” time by game players) or with “bars” that depict latency similar to those for mobile phone signal strength – i.e., more bars is a better, lower-latency connection.

Figure 5 depicts an example of latency exposure. There are two clients, Client A and Client B, each connected to the Server in the middle. Client A has a round-trip time of 25 milliseconds and Client B has a round-trip time of 100 milliseconds. The clients’ displays expose their round-trip time latencies in the upper corner of the screen, noting this as a “ping” value which is term commonly understood and used by network gamers.




Fig. 5. An example of latency exposure.

Vaghi et al. [139] describe several mechanisms by which latency can be revealed to players: a gauge showing an amount of latency, an area of uncertainty (larger with more latency), “shadows” and “ghosts” to represent the estimated position of objects based on latency, and visual emphasis indicating when events occur in time in order to make the presence of latency known.

Similarly, Fraser et al. [51] describe how the uncertainty due to latency can be rendered as a wireframe volume around an object, with the volume growing larger with an increase in latency, similar to the area of uncertainty.

Wikstrand et al. [141] use a “shadow” to represent the estimated position of a game avatar on the server based on the latency, similar to the shadows proposed by Vaghi et al. [139].

Gutwin et al. [58], propose magnitude “decorators” that are visual representations for enhancing a user’s awareness of latency. Decorator examples include numbers, translucency, color and even pulsing or oscillation to indicate latency.

4.2 Prediction

Prediction estimates the game state at a client without having the official game state from the server. Prediction takes advantage of the fact that the game client has information on the game world (e.g., object locations, world terrain) and processing capabilities (e.g., ability to render the game world), providing the player with world representations before they are confirmed by the server. The drawback with predictive techniques is that they can be wrong in that they may not accurately represent the game world as determined by the authoritative server. In such cases, there is inconsistency between the client game state and the server game state so that when the client game state is inevitably fixed to align with the server game state, the fixed is noticeable and may even be jarring to the player.

4.2.1 Self-Prediction predicts game state based on player input, but before getting confirmation from the server. Self-prediction is a natural technique for game programmers since most clients run a fully-functional game engine able to incorporate player input and compute game object interactions (e.g., physics, including collisions) without the server, and doing so can provide immediate feedback for the player.

Self-prediction is often called *client-side prediction* in papers, online blogs and player posts.

Figure 6 depicts an example of self-prediction. At time t_0 on the left, the player has a view of the game world that is consistent with the server’s (not shown). At time t_1 , the player has provided some game input (e.g., press the right arrow key) in order to move the green avatar to the right. The client assumes that this movement will be allowed by the server and so renders the world with the green avatar in the predicted location. Once the server receives and then responds to the player input at time t_2 there are two possibilities: in the first case, shown at the top, the server has accepted the input and the green avatar’s new position




Fig. 6. An example of self-prediction.

is confirmed; in the second case, shown at the bottom, the server has rejected the player input (e.g., the avatar is blocked by another object unknown to the client) and the client renders the world as specified by the server.

Burgess, Hanna, Shelly and Katchabaw [16, 74, 124] describe the use of a software design pattern based on the notion of “optimism” where a prediction on the client is assumed to be valid until the server does the official computation and returns the results. They provide additional techniques for synchronization and consistency checking.

Chen et al. [23] use a form of self-prediction called an “echo” to immediately show the player the effects of an action, even if additional delay is incurred before the official confirmation.

Wu and Ouhyoung [142] study “look-ahead” algorithms for 3d, head-mounted displays that predict object position and orientation. They study algorithms with different prediction complexities – simple to more complex. Also for a virtual environment with a head-mounted display, Tumanov et al. [137] describe predictions of “poses” for player based on the position and orientation in the motion space.

Brun et al. [15] mention self-prediction as it affects game state consistency in multiplayer games. They provide specific context for fairness in multi-player games, where players may have different amounts of latency.

Le et al. [80] capture movements of the hand in a touch interface, use these movements in a neural network, and predict the location of future touch positions. Similarly, Antoine et al. [6] use high-frequency data gathered from a computer mouse to predict the future velocity and position.

4.2.2 Other-prediction predicts game state for objects controlled by other players without having the actual state information from the server.

4.2.3 Interpolation predicts past states for objects controlled by other players based on the current state and previously known states. For example, the position of a vehicle can be interpolated to be in-between a past known location and its current known location.

Figure 7 depicts an example of interpolation. At the left at time t_0 , the client displays the known state of the game world at that time. At the right at time t_2 , the client has gotten an update on the world state. At




Fig. 7. An example of interpolation.

time t_1 in the middle, the client interpolates the position of the green avatar based on the known position in the past at t_0 and the known position in the future at t_2 . The client then renders this predicted state. Interpolation is typically used to visually smooth out game states rendered on the client in cases where the visual update rate happens more often than do updates received from the server.

Lee and Chang [86] evaluate how interpolation in the commercial first-person shooter game *Counter-Strike: Global Offensive* (Valve, 2012) improves player accuracy.

Sharkey et al. [122] describe interpolation via *local perception filters* that provide for a continuous view of the game world (i.e., without any abrupt transitions) as client game states are updated with remote user actions. As Smed et al. [128] describe, interpolation via local perception filters can be tuned to keep predictions closer to actual game state for local players than for remote players in order to keep a smooth view of the game state while mitigating the effects of latency.

Savery et al. [119] provide a toolkit for game developers to implement interpolation, including local perception filters and also smooth corrections to interpolate rendered game state towards the actual game state when a client gets updates. The authors also evaluate how smooth corrections are perceived by players compared to abrupt corrections when fixing inconsistent game state [117].

4.2.4 Extrapolation predicts future states for objects controlled by other players assuming current behaviors continue. For example, the position of a vehicle can be extrapolated to a future location based on its last known position and kinematic state (velocity, acceleration, orientation, and angular velocity).




Fig. 8. An example of extrapolation.

Figure 8 depicts an example of extrapolation. At the left at time t_0 , the client display renders the known state of the game world at that time. In the middle at time t_1 , the state of the world has been updated by the server and the client displays the green avatar's new position, moving to the right. At the right at time

t_2 , the client has not (yet) gotten an update on the world state, but can extrapolate the position of the green avatar based on the last known position at time t_1 and the avatar’s velocity. The client then renders this predicted state.

The information and processing used in the extrapolated prediction can vary from simple to complex, depicted in Figure 9 as a single dimension. Extrapolation techniques on the left use basic physics to make a prediction, the simplest of which is the last known position and the velocity. Small increments in processing and information are required for acceleration and orientation, with a bit more for forces (e.g., friction) and mass. In the middle, extrapolation techniques might use higher order information such as environmental factors, routes available for travel, or the intended destination of the predicted object. On the right are extrapolation techniques that consider even more complex information, such as past behaviors, and complex processing such as machine learning and full AI control of the predicted objects.




Fig. 9. Extrapolation information and/or processing dimension.

Extrapolation is by far the most widely-researched latency compensation technique in the literature.

Vaghi et al. [139] identify prediction as a way to anticipate future states (e.g., the “real” position of a ball) based on latency. Diot and Gauiter [38] implement a simple extrapolation technique for the open source game MiMaze [37] by replaying the position of a remote object if an update is not available when rendering. Mauve [98] describes how with extrapolation, mis-predictions may yield unexpected game states (e.g., dead players shooting).

Wikstrand et al. [141] describe extrapolation as one of several “predictor displays,” stating they work best when the extrapolated positions of objects are predictable. In their pong game implementation, they depict the extrapolation locations of paddles as shadows.

Shelly and Katchabaw [124] extrapolate the positions of server-controlled objects (e.g., an opponent) as part of their “optimistic” software design patterns to compensate for latency. Similarly, Jian et al. [71] describe extrapolation as an “optimistic technique” where the client computes and renders the location of server-controlled objects without waiting for the latency from the server response.

Many multiplayer-network games use the well-known extrapolation technique called *dead reckoning*¹ whereby object predictions are done with basic kinematic physics such as the last known position and velocity. McCarty et al. [100] describe using dead reckoning for Distributed Interactive Simulations (DIS) – an IEEE standard for multi-player combat simulation – for a virtual flight simulator. Mauve [97] describe dead reckoning use in DIS, including shortcomings to state consistency.

¹https://en.wikipedia.org/wiki/Dead_reckoning

Smed et al [127] describe dead reckoning as one of several aspects of networking in online games both for reduction in bitrates and in mitigating latency via extrapolation. Pantel and Wolf [106] propose extrapolations of two types: 1) positions of game objects via dead reckoning, and 2) positions of an input device (e.g., a joystick) which, in their racing game, results in moving a car to a predicted position.

Yu and Choy [145] and Zang and Georganas [147] propose adding an orientation threshold for sending dead reckoning updates above and beyond the position threshold typically used. Hanawa and Yonekura [61] refine dead reckoning to reduce prediction errors with additional mathematics (e.g., a Taylor expansion). Cai et al. [17] extend dead reckoning by adapting the thresholds used to determine when to send updates based on the area of interest and sensitivity. Kharitonov [75] proposes, and Almeida and Felinto [32] evaluate a dead reckoning extension that uses the motion patterns of the objects in addition to simple kinematic physics.

Aggarwal et al. [1] analyze use of dead reckoning with synchronized clocks and time stamps to improve extrapolation accuracy after updates are received. The authors extend their ideas to fairness for players with different latencies by equalizing errors from incorrect extrapolations [2]. Zhang et al. [148] propose using Aggarwal et al.’s dead reckoning with synchronized clocks combined with incoming delay (see Section 4.3.1.2) to further reduce prediction errors. Ishibashi et al. [62, 65, 78] also use extrapolation combined with incoming delay, providing a predicted position for updates outside of the incoming delay buffer. Jaya et al. [69] adjust the dead reckoning update threshold based on an interest management profile – the closer an entity is, the lower the threshold and vice versa.

Duncan and Gracanin [41] propose reducing possible state inconsistency in dead reckoning by pre-computing predicted states and sending updates *before* any inconsistency thresholds are breached. Palant et al. [104] explore extrapolation variants of dead reckoning, including using clock synchronization and a convergence algorithm that smooths out inconsistency in game states from extrapolation errors. Brun et al. [15] consider dead reckoning in terms of how it affects game state inconsistency for multi-player games with different latencies for each player. Roberts et al. [113] propose an extension to dead reckoning whereupon state updates are sent using not just a threshold for spatial inconsistency, but an accumulation of this inconsistency over time. Savery et al. [118–120], describe extrapolation, part of prediction, as a general technique to compensate for latency, with dead reckoning as a specific version. The authors provide a state consistency-centric view of the game world, including requirements on state divergence and different approaches to correct inconsistent states (e.g., smoothly or abruptly).

Li and Chen [87] extend dead reckoning by extrapolating object positions based on attraction (inferring that some objects may want to move closer to another) and repulsion (the opposite). Li et al. [88, 89] add the notion of personal habits, where the prediction takes into account actions that the predicted object would likely do based on past behavior. Similarly, Schirra [121] describes using the content of game objects in order to do dead reckoning using more than just position and velocity. For example, content-based dead reckoning might use “running along a path” or “playing a wall-pass to a team-mate” in order to extrapolate positions. Along these lines, Yahyavi et al. [144] consider the physics of a player avatar in predicting a location and the player’s interest in their surrounding environment in terms of how it affects movement. Similarly, Gao et al. [54] describe modeling a player’s behavior and using this behavior profile to extrapolate to positions of player-controlled objects. Their technique is combined with dead reckoning in an attempt to better predict actual avatar location in a pong game. And Chen and Liu [24] extend traditional dead reckoning

to an extrapolation technique that also uses human behavior and the virtual environment factors (e.g., the terrain) in doing predictions.

Lee and Chang [86] describe (and evaluate) what they call *lag compensation*, which is effectively extrapolation and time warp (see Section 4.3.1.1) combined.

Duarte et al. [40] propose using machine learning to infer if dead reckoning is able to accurately extrapolate positions – if not, another machine learning algorithm is used to make the predictions.

Cronin et al. [30] examine how prediction techniques (e.g., dead reckoning) may be incompatible with cheat-detection algorithms since extrapolation may be exploited in some cases.

4.2.5 Speculative Execution computes the game world state based on possible player input before it has actually happened, and adopts this pre-computed state if/when the input is provided. This enables the local game client to respond to player input immediately, matching the current game world to the pre-computed world, without waiting for the round-trip time to the server.




Fig. 10. An example of speculative execution.

Figure 10 depicts an example of speculative execution. On the left at time t_0 , the game client depicts the current game world state. At the top of the figure, the server has pre-computed and provide two possible game states that could exist at time t_1 , names State A and State B. These two states are based on possible input that the player could provide to the game at time t_0 – in this example, the player can either move the green avatar to the top right of the game world (resulting in State A) or to the bottom left of the game world (resulting in State B). When the actual input is provided, the state that corresponds to that input, either State A or State B, can be used immediately to represent the current game world state without the additional latency from client to server. The unused state is discarded.

Lee et al. [83] provide the only example of speculative executing found in the literature. Their *Outatime* system predicts future states based on a model of past user input and computes parallel states for input that is hard to predict. For cases of mis-prediction, Outatime uses corrective rendering to visually repair the mis-predictions, and time warp with state check-pointing to limit error propagation.

While not from a peer reviewed forum (therefore, not part of our survey), Kopietz et al. [77] from id Software (*Doom* and other titles) patented a speculative execution technique – rather than sending game states based on every possible combination of inputs, each input gets associated with a motion vector

describing how the rendered game frame would change in response to that input. For some games, when there is multiple inputs, this may allow adding up vectors and applying the sum to a rendered image, short-circuiting latency to the server.

4.3 Time Manipulation

Time Manipulation alters the virtual time (i.e., the time in the game world) for computing the game state and/or resolving player actions.

4.3.1 Time Delay buffers game state updates in order to provide for more uniform latency across clients. For example, delaying updates by $D - l_i$, where l_i is the latency for each client and D is the $\max(l_i)$, provides an equal latency for all players. Equal latencies across players can make the game more fair and game states more consistent.

4.3.1.1 Time Warp rolls back game state to when the player action occurred on the client, applies the action, then rolls the game state forward to the current time.

Figure 11 depicts an example of time warp. The figure shows the game world for a shooter game on a Client and the Server, with time advancing from top to bottom. The player on the client is shooting at a green avatar that is moving right to left, with the “plus” sign in the middle representing a weapon reticle. At time t_0 at the client, the green avatar is to the right of the reticle, moving into the reticle at time t_1 where the player pulls the trigger, and that action is sent to the server arriving just after time t_2 . Meanwhile, on the server, the green avatar moves past the reticle at time t_1 and has continued right at time t_2 . When the action arrives at the server, the server “warps” time back to when the action occurred at the client at time t_1 , applying the action to the world representation at that time.




Fig. 11. An example of time warp.




Fig. 12. An example of “shot around the corner”.

However, resolving an event in the past and rolling it forward may cause already rendered client game states to be inconsistent with the new game state. This is a well-known artifact of some shooting games, and can result in “shot around the corner” as it’s known by players, shown in Figure 12. At time t_1 , at the blue avatar’s client, the green avatar is in sight and the blue player fires. However, by time t_2 when

the server receives the update, the green avatar has reached a safe position around the corner from the blue player and cannot be targeted. However, with time warp, the server, upon receiving the blue player's action at time t_2 , rolls back time to the green avatar's position at time t_1 and applies the action. This hits the green avatar. Rolling the game world forward, with the green avatar hurt or killed, may feel like being “shot around the corner” for the green player.

Time warp is often called, somewhat confusingly, *latency compensation* or *lag compensation* in some papers and, more often, in online blogs and player posts. It is also sometimes simply called *rollback*. Also somewhat confusing, some image warping techniques (e.g., [43]) are called *time warp* whereas we classify them as latency concealment (see Section 4.1.1).

Jefferson [70] first proposed *virtual time* as a paradigm for distributed computation, with time warp as an implementation. Although multi-player games were not identified as a use at that time, distributed discrete event simulations were.

Mauve [98] describes how timestamps and time warp can be used to overcome mistakes made using extrapolation. In particular, an extrapolated state update may miss a key event, such as a player being killed that time warp can roll back to correct. Mauve [97] and Mauve et al. [99] provide a formalization of time warp in the context of continuous, interactive media, such as computer games. Jiang et al. [71] mention time warp in their survey as a means to overcome inconsistencies caused by prediction.

Savery et al. [120] mention the use of time warp and the “shot around the corner” problem in commercial first-person shooter games, illustrated Figure 12 whereby time warp can undo a players move to a safe location to being damaged. Lee and Chang [84, 86] describe and evaluate the “shot around the corner” problem. Subsequently, they [85] describe how commercial first-person shooters provide a limit on how far back a server rolls back time and propose an advanced time warp technique to prevent “shot around the corner” whereby a client can identify the player is currently safe (using their local time) and prevent rollback.

Sun and Claypool [135] implement and evaluate time warp for a cloud-based game streaming system.

4.3.1.2 Incoming time delay buffers player actions before applying them so that actions arrive (and are applied) at all clients simultaneously.




Fig. 13. An example of incoming delay.

Figure 13 depicts an example of incoming delay. The figure depicts downward timelines on a server and 2 clients, Client A and Client B. At time t_0 , player A responds to the game with an action. That action is

immediately sent to the Server and then to Client B. Because of the latency for Client A to the Server and the Server to Client B, that action does not arrive at Client B until time t_1 . In order to have the action executed simultaneously on both clients, Client A adds an incoming delay equal to the latency from Client A to the Server plus the latency from the Server to Client B before applying the action so that Client A also applies the action at time t_1 . Note, while the extra delay is shown at the client in this example, it can also be used at the server.

Mauve et al. [97, 99] apply and formalize incoming delay as depicted in Figure 13. Savery and Graham [119] provide a toolkit for implementing incoming delay.

Diot et al. [38] describe a bucket synchronization mechanism whereby all game state updates received by a client are stored (delayed) until their corresponding time interval.

Lin et al. [91] use incoming delay on both clients and server in order to synchronize game states temporally, calling their technique *sync-in* and *sync-out*.

Paik et al. [103] apply incoming delay at the server to account for different client latencies, with the delay time adjusted based on the number of clients within virtual proximity to each other.

Savery et al. [118] use incoming delay in order to improve consistency. Brun et al. [13] propose tuning the amount of incoming delay to find the best tradeoff between consistency and responsiveness, also using different amounts of incoming delay for different game actions. Chen et al. [23] combine self-prediction (see Section 4.2.1) with incoming delay – the incoming delay reduces state inconsistency across players and the self-prediction alleviates some of the reduced responsiveness caused by the extra delay. Zhang et al. [148] combine incoming delay with extrapolation (see Section 4.2.4) – incoming delay to reduce inconsistency before updates are received and extrapolation to reduce inconsistency after updates are sent. Stuckel and Gutwin [134] evaluate incoming delay as well as incoming delay plus the prediction “echo” by Chen et al. [23]. Ishibashi et al. [62, 65, 78] use a standard incoming delay technique that buffers incoming game state updates up to the buffer period, then predicts (extrapolates) object positions that do not have an update.

Brun et al. [15] mention using incoming delay as a way to adjust fairness among all players (i.e., players closer to the server will be delayed to the same latency as players further away), which has a tradeoff of decreasing playability for some players. Similarly, Le and Liu [79] use incoming delay on the server by holding those packets from clients with lower than average delay so as to equalize latency across all clients.

Savery et al. [120] propose incoming delay in two forms: 1) immediately send out a player’s input, but delay applying it to reduce inconsistency across players (they call this *local lag*), and 2) an analogy to a streaming media applications whereby state arriving from a game server is not immediately applied but rather buffered to smooth out playback (they call this *remote lag*). Xu and Wah [143] describe an incoming delay approach as well as a hybrid scheme combining incoming delay and prediction – in this case interpolation to smooth out updates. Zhao et al. [149] use incoming delay to smooth input in a system supporting legacy games.

4.3.1.3 Outgoing time delay buffers remote game state updates before sending to provide equal latencies across clients. Outgoing delay is similar to incoming delay, but while incoming delay adds the delay after receiving a message (e.g., a game action), outgoing delay adds the delay *before* sending a message.




Fig. 14. An example of outgoing delay.

Figure 14 depicts an example of outgoing delay. The figure depicts timelines advancing top to bottom on a server and 2 clients, with Client A having a lower latency to the server than Client B. At time t_0 , the server has a game world update message to send to both Client A and Client B and sends the message to Client B. However, since Client A has a lower latency than Client B, the server delays sending the message until time t_1 . Setting the outgoing delay ($t_1 - t_0$) to the difference in latency for Client B and Client A means the message arrives at Client A and Client B simultaneously, at time t_2 .

Zander et al. [146] implement outgoing delay on a game server to add delay to outgoing messages to each player so as to equalize their latencies. The maximum total target delay threshold is adjusted depending upon game type.

Brun et al. [15] propose using outgoing delay in servers so as to manage fairness across clients with different latencies.

Wikstrand et al. [141] describe an “input buffering” scheme which uses outgoing delay to add latency to outgoing messages to other players so as to equalize their arrival times (they consider a peer-to-peer architecture).

Kaiser et al. [73] implement outgoing delay by concatenating game update messages at the client into one larger packet for an intended delay period before sending to the server. The server does something similar for updates to the clients.

Li et al. [90] implement outgoing delay for a cloud-based game platform where the amount delayed is inversely proportional to the latency for each thin client. They propose a maximum threshold to avoid extremely high latencies that might be required to accommodate some clients.

4.4 World Adjustment

World Adjustment modifies game states to decrease difficulty akin to configurations with lower latency.

4.4.1 Control Assistance adjusts the outcome of player input to accommodate for inaccuracies due to latency.

Figure 15 depicts an example of one type of control assistance – “target magnetism”. The blue avatar aims and shoots at the green avatar. Normally, the shot would miss, as indicated by the solid red line. However, with the target magnetism control assistance technique, the bullet trajectory is altered to bend towards




Fig. 15. An example of control assistance.

the green avatar, making it easier for the blue player to hit the target. Other control assistance technique examples include “sticky targets” that reduce the cursor gain while near a target and “aim dragging” that cause the cursor to follow the direction of an intended target. The amount of assistance provided can be set proportional to the latency.

Mandryk and Gutwin [96] test the utility and perceptibility of the sticky targets control assistance technique applied to a computer mouse, albeit without additional latency.

Bateman et al. [8] compare three control assistance techniques for mouse pointing: cursors that cover an area, “gravity” towards a target, and the aforementioned sticky targets.

Vicencio-Moriera et al. [140] examine control assistance for aiming with a mouse in a first-person shooter game considering 5 techniques: automatically locking on a target, bullet magnetism as in Figure 15, and the aforementioned area cursor, sticky targets, and target gravity. Again, their focus is on ability to help the player, not necessarily to overcome latency.

Ivkovic et al. [66] study the ability of control assistance for aiming a computer mouse in a first-person shooter game and tracking a target. Their test conditions were for local system latency not network latency, but pertain to cloud-based game streaming with network latency. Sabet et al. [115] also examine control assistance for cloud-based game streaming – in their case, for area cursors in 2d games with mouse-based aiming.

4.4.2 Attribute Scaling increases or decreases numeric attributes of objects and other game world parameters to adjust game difficulty so as to make player actions easier to complete with higher latency. For example, decreasing the size of obstacles can make them easier to avoid during navigation or increasing the speed of the player’s avatar can make it more nimble. The intent can also be to make the action more resilient to mistakes. For example, avatar durability can be increased with extra lives or increasing the health or armor attributes.

Figure 16 depicts an example of attribute scaling for a game in which a player tries to navigate the green avatar from left to right through narrow gaps between obstacles. There are two different latency scenarios depicted for the same game condition. On the left is a low latency scenario which has relatively narrow gaps between the obstacles since the player’s control is not significantly impacted by the latency. In contrast, on the right is a high latency scenario in which the game attributes have been scaled to make the gaps between the obstacles larger since the player’s control is more sluggish with the latency.

Vaghi et al. [139] mention adapting the pace of interaction required by a game with latency (i.e., the user interaction rate would slow down with higher latency).




Fig. 16. An example of attribute scaling.

Xu and Wah [143] extend the hitting time in a fighting game based on latency, where higher latency players would have a longer time period to execute a strike.

Shen and Zhou [125] change the speed of the ball in a pong game in relation to latency and the object's proximity to a player-controlled avatar that may interact with it – i.e., the ball slows down when nearing a player controlled object when there is high latency.

Sabet et al. [114] evaluate adjusting game attributes in response to higher latency by: 1) increasing the deadline for a game action (i.e., giving the player more time to execute an action), and 2) decreasing the precision required for a game actions. The authors [115] extend this work with additional attribute scaling for: 3) predictability, where randomness is decreased with latency, 4) number of required actions (NRA) per minute, where NRA decreases with latency, and 5) consequence where the impact of a detrimental action (e.g., a collision) is decreased with higher latency.

Lee et al. [82] propose adjusting the geometry of a game world based on latency. Specifically, they adjust the gap distance between pipes based on latency for the game Flappy Bird. Salay and Claypool [116] evaluate the same – pipe gap distance based on latency for Flappy Bird – in addition to allowing the player to manually adjust game attributes in response to latency.

4.5 Summary

Table 2 provides a summary table of the taxonomy in Figure 3 with the techniques rotated and presented top to bottom. For each technique, the corresponding publications are shown on the right.

From the table, extrapolation is the most published latency compensation technique, with 41 papers (half) in this area, more than twice that of the next closest – incoming time delay with 19. Correspondingly, other-prediction and time delay are the most published techniques, belonging to the general group of prediction and time manipulation, respectively.

5 Latency Compensation in Commercial Games

Most commercial games are “black boxes” wherein it is difficult to ascertain what, if any, latency compensation techniques are being deployed. However, the developers for some games have provided specific mention and even details about latency compensation techniques used in their software. This section provides an overview of the techniques and the games found in these developer blogs.

Table 2. Publications for each latency compensation technique. Color intensity is proportional to number of publications.

Technique			Publication	
Latency Compensation	Feedback	Concealment	[12, 16, 43, 74, 74, 76, 124, 129]	
		Exposure	[51, 58, 139, 141]	
	Prediction	Self-prediction	[6, 15, 16, 23, 74, 80, 124, 137, 142]	
		Other-prediction	Interpolation	[86, 117, 119, 122, 128]
			Extrapolation	[1, 2, 15, 17, 24, 30, 32, 37, 38, 40, 41, 61]
				[54, 62, 65, 69, 71, 75, 78, 86–89, 97]
				[98, 100, 104, 106, 113, 118–121, 124]
				[127, 139, 141, 144, 145, 147, 148]
		Speculative Execution	[83]	
	Time Manipulation	Time Delay	Incoming	[13, 15, 23, 38, 62, 65, 78, 79, 91, 97]
			Outgoing	[99, 103, 118–120, 134, 143, 148, 149]
		Time Warp	[70, 71, 84–86, 97–99, 120, 135]	
	World Adjustment	Control Assistance	[8, 66, 96, 115, 140]	
		Attribute Scaling	[82, 114–116, 125, 139, 143]	

The games listed are not meant to indicate that these are the only games that use latency compensation techniques, far from it – it is our expectation and experience that most if not all AAA network games use some form of latency compensation. Instead, this section includes known games where the game developers themselves describe through either Web documents or online presentations the latency compensation techniques implemented in their games.

In locating these blogs, we started with a list of 3 or 4 blogs we knew about previously.

We then searched two specific online forums that are known for having technical content posted by game developers themselves: the *Game Developer’s Conference*² and *Gamasutra*.³

Next, we visited the forums that are maintained by the publishers for the top-10 esports games [111] (the list is based on viewership and prize pools) and searched for latency compensation posts made by game developers. The list of games with forum citations are shown in Table 3, given in list order.

Lastly, we did general Internet searches for latency compensation blogs – Web pages and videos – including only those made by game developers about their own commercial games.

Table 4 summarizes the results. *Year* is when the paper was published. *Compensation Technique* maps the latency compensation technique(s) in the blog to our taxonomy. Many of the blogs indicate the games implement more than one technique. *Game* provides the names of the game(s) described in the blog. *Genre* provides for the general class of games to which the targeted games belong: “FPS” (first-person shooter), “RPG” (role playing game), “Fighting”, and “Sports”. *Range* indicates the latency values mentioned in the blog, or “unspecified” if none.

From the table, extrapolation (10) is referenced the most, followed by time warp (9), incoming delay (7), interpolation (5), self-prediction (4) and latency exposure (2). Two-thirds (12 of 20) of the references are to FPS games, probably owing to their sensitivity to latency [26] and heavy use in competitive gaming (e.g., esports).

²<https://gdconf.com/>

³<https://www.gamasutra.com/>

Table 3. Top 10 esports games searched for latency compensation posts.

Game	Publisher	Year	Forum
League of Legends	Riot Games	2009	[81]
CS:GO	Valve Corporation	2012	[31]
Fortnite	Epic Games	2017	[42]
DOTA2	Valve Corporation	2013	[39]
Hearthstone	Blizzard Entertainment	2014	[48]
Call of Duty	Activision	2003	[27]
Overwatch	Blizzard Entertainment	2016	[49]
Rainbow 6 Siege	Ubisoft	2015	[50]
Rocket League	Psyonix	2015	[132]
NBA 2K	Sega Sports	1999	[47]

Table 4. Developer blogs with details on latency compensation implementations for commercial network games.

Year	Link	Compensation Technique	Game	Genre	Range
1996	[19]	Self-Prediction	QuakeWorld	FPS	up to 200 ms
2001	[11]	Extrapolation	Age of Empires	RTS	unspecified
2001	[10]	Extrapolation, Interpolation, Time Warp	Half-life	FPS	unspecified
2011	[3]	Extrapolation, Incoming Delay, Latency Exposure	Halo	FPS	up to 300 ms
2012	[18]	Self-Prediction, Extrapolation, Time Warp	GGPO (library)	Fighting	unspecified
2016	[56]	Interpolation, Self-Prediction	Call of Duty: Black Ops III	FPS	unspecified
2016	[46]	Extrapolation, Incoming Delay, Interpolation, Self-Prediction, Time Warp	Overwatch	FPS	unspecified
2016	[136]	Extrapolation, Time Warp	MechWarrior Online	FPS	unspecified
2017	[34]	Extrapolation, Interpolation	Watch Dogs 2	RPG	unspecified
2017	[28]	Self-Prediction, Time Warp	Half-life	FPS	up to 200 ms
2018	[33]	Extrapolation	Rocket League	Sports	up to 600 ms
2018	[29]	Extrapolation	Rocket League	Sports	unspecified
2018	[131]	Incoming Delay, Time Warp	Injustice 2, Mortal Kombat	Fighting	up to 300 ms
2018	[5]	Interpolation, Time Warp	Unity's FPS Sample Game	FPS	up to 200 ms
2019	[60]	Extrapolation, Incoming Delay	COD: Modern Warfare	FPS	unspecified
2019	[108]	Incoming Delay, Time Warp	Various Fighting games	Fighting	up to 150 ms
2019	[45]	Extrapolation, Incoming Delay	Overwatch	FPS	up to 1000 ms
2020	[133]	Time Warp	Valorant	FPS	unspecified
2020	[112]	Time Warp	Valorant	FPS	unspecified
2020	[35]	Incoming Delay, Latency Exposure, Self-Prediction	Valorant	FPS	20 - 50 ms

6 Discussion

Based on the sheer number of publications, there has been the most research done in latency compensation via prediction, specifically extrapolation. However, while prediction is well-researched in general, the area of

speculative execution is not. Improvements to speculative execution may rely upon emerging AI techniques to predict player input and, hence, game state before it happens. While predicting player actions in general may be difficult, doing so for the constrained conditions provided by many games seems possible.

Manipulating the virtual time in game worlds has been fruitful, most prominently by delaying incoming actions to smooth input and provide fairness, and in “warping” time to resolve past actions. Commercial games make heavy use of both of these techniques.

Conversely, there is substantially less research in feedback, despite feedback perhaps being the most direct connection to the player via response to input. Latency exposure in the form of “ping” values, is common to many games, but research on such exposures mitigate latency as well as more advanced exposure approaches merit additional exploration.

There is considerable potential in research in world adjustments, too, with most control assistance techniques not having been assessed for latency and approaches to attribute scaling as broad and deep as games themselves.

In general, many latency compensation techniques could use additional evaluation, both to study a broader range of latencies but also with a broader set of users. Ideally, a representative sample from a user study would cover the broad demographic range of today’s gamers.

While the blogs obtained from game developers provide evidence of latency compensation in commercial games, the extent to which academic research has been incorporated into commercial products is not well known. In fact, the relatively narrow range of techniques in Table 4 suggests many techniques proposed by the scientific literature may not yet be incorporated into commercial games.

Latency poses a particular challenge for a relatively recent game system type – that of cloud-based game streaming,⁴ such as Microsoft’s *Xcloud*, Amazon’s *Luna*, NVidia’s *GeForce Now*, Google’s *Stadia* and Sony’s *PlayStation Now*. Unlike in traditional network games, with cloud-based game streaming the client does not have the game state nor does it typically do any “heavyweight” processing such as 3d graphics rendering. Instead, the game frames are rendered at the cloud-based server and streamed down to the client similar to video. The client captures player input (e.g., mouse movements, button presses) and sends those back to the server for processing. In this type of system, the “thin” client cannot do the processing required by some of the latency compensation techniques. Specifically, the green shaded techniques in Figure 3 labeled “Client” cannot be done in cloud-based game streaming. Cloud-based game streaming systems generally can only apply the “Server” or “Either” techniques. This suggests the opportunity and challenge to develop new techniques to compensate for latency in cloud-based game streaming.

7 Summary

Computer games continue to grow in popularity for entertainment and professionally through esports. Computer games are also often multi-player, connecting computers over a network to enable geographically separated players to interact with each other in the same virtual world. These multi-player, network games need to exchange player actions and game state across the network fast and frequently to provide for a smooth, interactive player experience. Unfortunately, network latency adds delay for game updates from

⁴https://en.wikipedia.org/wiki/Cloud_gaming

the server and player actions from the client, degrading the quality of experience by making the game less responsive and increasing inconsistency across client game states.

Latency compensation techniques use software at the client or server (or both) to mitigate the negative affects of network latency. While some latency compensation techniques having been established decades ago and well-used in the commercial game industry today, to the best of our knowledge, the space of latency compensation techniques has not been surveyed, nor have techniques been grouped based on their characteristics. Such a survey and organization can provide guidance for game and game system developers for techniques to deploy and for researchers that seek to invent new techniques.

We located and surveyed 82 peer-reviewed publications that dealt with latency compensation techniques. Our records list the type of techniques used and summary of evaluation, including game(s) studied, type of evaluation and latency range. Our search also yielded evidence of latency compensation for commercial games, with 20 developer blogs detailing their latency compensation implementations. The latency compensation techniques are placed into a novel taxonomy based on similarity, with four main categories - feedback, prediction, time manipulation and world adjustment - leading to 11 technique types - concealment, exposure, self-prediction, interpolation, extrapolation, speculative execution, incoming delay, outgoing delay, time warp, control assistance and attribute scaling. The most popular latency compensation categories are prediction in the form of extrapolation and time manipulation in the form of incoming delay.

Future work is to continue to develop and evaluate novel latency compensation techniques for multiplayer, network computer games. In particular, there is a need for latency compensation techniques for cloud-based game streaming systems – since such systems cannot do heavyweight computation on the client, viable techniques must be server-side only. Future work is also to assess latency compensation techniques on the increasingly broad range of gaming devices and inputs, such as touch screens and virtual reality.

References

- [1] Sudhir Aggarwal, Hemant Banavar, Amit Khandelwal, Sarit Mukherjee, and Sampath Rangarajan. 2004. Accuracy in Dead-Reckoning Based Distributed Multi-Player Games. In *Proceedings of 3rd ACM Workshop on Network and System Support for Games (NetGames)*. Portland, OR, USA, 161–165.
- [2] Sudhir Aggarwal, Hemant Banavar, Sarit Mukherjee, and Sampath Rangarajan. 2005. Fairness in Dead-Reckoning Based Distributed Multi-Player Games. In *Proceedings of 4th ACM Workshop on Network and System Support for Games (NetGames)*. Association for Computing Machinery, Hawthorne, NY, USA, 1–10.
- [3] David Aldbridge. 2011. I Shot You First: Networking the Gameplay of Halo:Reach. YouTube. <https://www.youtube.com/watch?v=h47zZrqjgLc>.
- [4] Rahul Amin, France Jackson, Juan E. Gilbert, Jim Martin, and Terry Shaw. 2013. Assessing the Impact of Latency and Jitter on the Perceived Quality of Call of Duty Modern Warfare 2. In *Proceedings of HCI – Users and Contexts of Use*. Las Vegas, NV, USA, 97–106.
- [5] Peter Andreasen. 2018. Deep dive into networking for Unity’s FPS Sample game - Unite LA. YouTube. https://www.youtube.com/watch?app=desktop&v=k6JTaFE7SYI&ab_channel=Unity.
- [6] Axel Antoine, Sylvain Malacria, and G ry Casiez. 2018. Using High Frequency Accelerometer and Mouse to Compensate for End-to-End Latency in Indirect Interaction. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Association for Computing Machinery, Montreal, QC, Canada, 1–11. <https://doi.org/10.1145/3173574.3174183>
- [7] Grenville Armitage. 2003. An Experimental Estimation of Latency Sensitivity in Multiplayer Quake 3. In *Proceedings of the 11th IEEE International Conference on Networks (ICON)*. Sydney, Australia.
- [8] Scott Bateman, Regan L. Mandryk, Tadeusz Stach, and Carl Gutwin. 2011. Target Assistance for Subtly Balancing Competitive Play. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Vancouver, BC, Canada.

- [9] Tom Beigbeder, Rory Coughlan, Corey Lusher, John Plunkett, Emmanuel Agu, and Mark Claypool. 2004. The Effects of Loss and Latency on User Performance in Unreal Tournament 2003. In *Proceedings of ACM Network and System Support for Games Workshop (NetGames)*. Portland, OR, USA.
- [10] Yahn Bernier. 2001. Latency Compensating Methods in Client/Server In-game Protocol Design and Optimization. In *Proceedings of the Game Developers Conference (GDC)*. San Francisco, CA, USA. Online: <https://www.gamedevs.org/uploads/latency-compensation-in-client-server-protocols.pdf>.
- [11] Paul Bettner and Mark Terrano. 2001. 1500 Archers on a 28.8: Network Programming in Age of Empires and Beyond. Gamasutra. https://www.gamasutra.com/view/feature/131503/1500_archers_on_a_288_network_.php.
- [12] Ben Boudaoud, Pyarelal Knowles, Joohwan Kim, and Josef Spjut. 2021. Gaming at Warp Speed: Improving Aiming with Late Warp. In *Proceedings of the ACM SIGGRAPH Emerging Technologies Workshop*. Virtual.
- [13] J. Bran, F. Safaei, and P. Boustead. 2004. Tailoring Local Lag for Improved Playability in Wide Distributed Network Games. In *Proceedings of Australian Telecommunication Networks and Applications Conference*. Sydney, Australia.
- [14] Bob Briscoe, Anna Brunstrom, Andreas Petlund, David Hayes, David Ros, Ing-Jyh Tsang, Stein Gjessing, Gorrry Fairhurst, Carsten Griwodz, and Michael Welzl. 2016. Reducing Internet Latency: A Survey of Techniques and their Merits. *IEEE Communications Surveys & Tutorials* 18, 3 (2016).
- [15] Jeremy Brun, Farzad Safaei, and Paul Boustead. 2006. Managing Latency and Fairness in Networked Games. *Commun. ACM* 49, 11 (Nov. 2006), 46–51. <https://doi.org/10.1145/1167838.1167861>
- [16] Shayne Burgess and Michael Katchabaw. 2006. Design and Implementation of Optimistic Constructs for Latency Masking In Online Video Games. In *The 2nd annual North American Game-On Conference (GameOn'NA 2006)*. Monterey, CA, USA.
- [17] Wentong Cai, Francis B. S. Lee, and L. Chen. 1999. An Auto-Adaptive Dead Reckoning Algorithm for Distributed Interactive Simulation. In *Proceedings of the Thirteenth IEEE Workshop on Parallel and Distributed Simulation (PADS)*. Atlanta, GA, USA.
- [18] Tony Cannon. 2012. Fight the Lag! The Trick Behind GGPO's Low-latency Netcode. *Game Developer Magazine* (Sept. 2012). https://drive.google.com/file/d/1cV0fY8e_SC1hIFF5E1rT8XRVRzPjU8W9/view.
- [19] John Carmack. 1996. Here is the New Plan. Blog. <https://fabiansanglard.net/quakeSource/johnnc-log.aug.htm>.
- [20] Elie Cattán, Amélie Rochet-Capellan, Pascal Perrier, and François Berard. 2015. Reducing Latency with a Continuous Prediction: Effects on Users' Performance in Direct-Touch Target Acquisitions. In *Proceedings of the International Conference on Interactive Tabletops & Surfaces (ITS)*. Madeira, Portugal.
- [21] Nilanjan Chaudhuri, Swakshar Ray, Rajat Majumder, and Balarko Chaudhuri. 2009. A New Approach to Continuous Latency Compensation With Adaptive Phasor Power Oscillation Damping Controller (POD). *IEEE Transactions on Power Systems* 25, 2 (May 2009).
- [22] K. Chen, P. Haung, G. Wang, C. Huang, and C. Lee. 2006. On the Sensitivity of Online Game Playing Time to Network QoS. In *Proceedings of IEEE Infocom*. Barcelona, Spain.
- [23] Ling Chen, Gen-Cai Chen, Hong Chen, Jack March, Steve Benford, and Zhi-Geng Pan. 2007. An HCI Method to Improve the Human Performance Reduced by Local-lag Mechanism. *Interacting with Computers* 19, 2 (March 2007).
- [24] Youfu Chen and Elvis S. Liu. 2018. Comparing Dead Reckoning Algorithms for Distributed Car Simulations. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*. Rome, Italy.
- [25] Mark Claypool. 2005. The Effect of Latency on User Performance in Real-Time Strategy Games. *Elsevier Computer Networks, Special Issue on Networking Issues in Entertainment Computing* 49, 1 (Sept. 2005), 52–70.
- [26] Mark Claypool and Kajal Claypool. 2006. Latency and Player Actions in Online Games. *Commun. ACM* 49, 11 (Nov. 2006).
- [27] codforums.com. [n.d.]. COD Forum. <https://www.codforums.com/>.
- [28] Valve Developer Community. 2019. Source Multiplayer Networking. Valvesoftware. https://developer.valvesoftware.com/wiki/Source_Multiplayer_Networking.
- [29] Jared Cone. 2018. It IS Rocket Science! The Physics of Rocket League Detailed. YouTube. <https://youtu.be/ueEmiDM94IE?t=23m33s>.
- [30] Eric Cronin, Burton Filstrup, and Sugih Jamin. 2003. Cheat-proofing Dead Reckoned Multiplayer Games. In *Proceedings of the 2nd International Conference on Application and Development of Computer Games*. Hong Kong, China.
- [31] csgoforum.com. [n.d.]. CSGO Forum. <https://www.csgoforum.com/>.
- [32] Luis Fernando Kawabata de Almeida and Alan Salvany Felinto. 2018. Evaluation of the Motion-Aware Adaptive Dead Reckoning Technique Under Different Network Latencies Applied in Multiplayer Games. In *Proceedings of the 17th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*. Foz do Iguaçu, Brazil.

- [33] Rocket Science-Halfway Dead. 2018. RL Netcode and Lag explained - Rocket Science #16. YouTube. <https://www.youtube.com/watch?v=c373LsgIXBc>.
- [34] Matt Delbosc. 2017. Replicating Chaos: Vehicle Replication in Watch Dogs 2. YouTube. https://www.youtube.com/watch?v=_8A2gzRrWLk.
- [35] Matt deWet and David Straily. 2020. Peeking into Valorant's Netcode. Riot Games. <https://technology.riotgames.com/news/peeking-valorants-netcode>.
- [36] Matthias Dick, Oliver Wellnitz, and Lars Wolf. 2005. Analysis of Factors Affecting Players' Performance and Perception in Multiplayer Games. In *Proceedings of the ACM Workshop on Network and Systems Support for Games (NetGames)*. Hawthorn, NY, USA.
- [37] C. Diot and L. Gautier. 1998. Design and Evaluation of MiMaze - a Multi-player Game on the Internet. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*. Austin, TX, USA.
- [38] C. Diot and L. Gautier. 1999. A Distributed Architecture for Multiplayer Interactive Applications on the Internet. *Network Magazine of Global Internetworking* 13, 4 (July 1999), 6–15. <https://doi.org/10.1109/65.777437>
- [39] dota2.com. [n.d.]. dota2 Forum. <https://dev.dota2.com/>.
- [40] Elias P Duarte Jr, Aurora TR Pozo, and Pamela Beltrani. 2020. Smart Reckoning: Reducing the Traffic of Online Multiplayer Games Using Machine Learning for Movement Prediction. *Entertainment Computing* 33 (2020), 100336.
- [41] Thomas Duncan and Denis Gracanin. 2003. Algorithms and Analyses: Pre-reckoning Algorithm for Distributed Virtual Environments. In *Proceedings of the 35th Conference on Winter Simulation: Driving Innovation (WSC)*. New Orleans, LA, USA.
- [42] epicgames.com. [n.d.]. Fortnite Forum. <https://www.epicgames.com/fortnite/en-US/home>.
- [43] Daniel Evangelakos and Michael Mara. 2016. Extended TimeWarp Latency Compensation for Virtual Reality. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. Redmond, WA, USA.
- [44] Parinaz Farajiparvar, Hao Ying, and Abhilash Pandya. 2020. A Brief Survey of Telerobotic Time Delay Mitigation. *Frontiers in Robotics and AI* 7, 578805 (Dec. 2020).
- [45] Timothy Ford. 2019. Overwatch Gameplay Architecture and Netcode. YouTube. https://www.youtube.com/watch?app=desktop&v=W3aieHjyNvw&ab_channel=GDC.
- [46] Tim Ford and Philip Orwig. 2016. Fighting Latency on Call of Duty: Black Ops III. YouTube. <https://www.youtube.com/watch?v=vTH2ZPgYujQ>.
- [47] forums.2k.com. [n.d.]. nba2k Forum. <https://forums.2k.com/forumdisplay.php?214-NBA-2K>.
- [48] forums.blizzard.com. [n.d.]. heartstone Forum. <https://us.forums.blizzard.com/en/hearthstone/>.
- [49] forums.blizzard.com. [n.d.]. Overwatch Forum. <https://us.forums.blizzard.com/en/overwatch/>.
- [50] forums.ubisoft.com. [n.d.]. rs6 Forum. <https://forums.ubisoft.com/forumdisplay.php/64-Rainbow-Six-Siege>.
- [51] Mike Fraser, Tony Glover, Ivan Vaghi, Steve Benford, Chris Greenhalgh, Jon Hindmarsh, and Christian Heath. 2000. Revealing the Realities of Collaborative Virtual Reality. In *Proceedings of the Third International Conference on Collaborative Virtual Environments (CVE)*. San Francisco, CA, USA.
- [52] Sebastian Friston, Per Karlström, and Anthony Steed. 2016. The Effects of Low Latency on Pointing and Steering Tasks. *IEEE Transactions on Visualization and Computer Graphics* 22, 5 (May 2016), 1605–1615.
- [53] Tobias Fritsch, Hartmut Ritter, and Jochen H. Schiller. 2005. The Effect of Latency and Network Limitations on MMORPGs: a Field Study of Everquest 2. In *Proceedings of the 4th ACM Network and System Support for Games (NetGames)*. Hawthorne, NY, USA.
- [54] Chen Gao, Haifeng Shen, and M. Ali Babar. 2016. Concealing Jitter in Multi-player Online Games through Predictive Behaviour Modeling. In *Proceedings of the IEEE International Conference on Computer Supported Cooperative Work in Design (CSCWD)*. Nanchang, China.
- [55] Steven Gargolinski, Christopher St. Pierre, and Mark Claypool. 2005. Game Server Selection for Multiple Players. In *Proceedings of the 4th ACM Network and System Support for Games (NetGames)*. Hawthorne, NY, USA.
- [56] Benjamin Goyette. 2016. Fighting Latency on Call of Duty: Black Ops III. YouTube. <https://www.youtube.com/watch?v=EtLHLfNpu84>.
- [57] Grand View Research. 2020. Video Game Market Size, Share & Trends Report. Report ID: GVR-4-68038-527-4. <https://tinyurl.com/y4s5vfrm>.
- [58] Carl Gutwin, Steve Benford, Jeff Dyck, Mike Fraser, Ivan Vaghi, and Chris Greenhalgh. 2004. Revealing Delay in Collaborative Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Vienna, Austria.
- [59] Carl Gutwin, Jeff Dyck, and Jennifer Burkitt. 2003. Using Cursor Prediction to Smooth Telepointer Jitter. In *Proceedings of the ACM SIGGROUP Conference on Supporting Group Work*. Sanibel Island, FL, USA, 294–301.

- [60] Paul Haile and Mitch Sanborn. 2019. Call of Duty: Modern Warfare Netcode Explained! YouTube. <https://www.youtube.com/watch?v=tCpYV4kIzE>.
- [61] Dai Hanawa and Tatsuhiko Yonekura. 2006. A Proposal of Dead Reckoning Protocol in Distributed Virtual Environment based on the Taylor Expansion. In *Proceedings of the International Conference on Cyberworlds*. Lausanne, Switzerland, 107 – 114. <https://doi.org/10.1109/CW.2006.10>
- [62] Yusuke Hara, Yutaka Ishibashi, Norishige Fukushima, and Shinji Sugawara. 2012. Adaptive Delta-Causality Control Scheme with Dynamic Control of Prediction Time in Networked Haptic Game. In *Proceedings of the ACM Workshop on Network and Systems Support for Games (NetGames)*. Venice, Italy.
- [63] O. Hohlfeld, H. Fiedler, E. Pujol, and D. Guse. 2016. Insensitivity to Network Delay: Minecraft Gaming Experience of Casual Gamers. In *Proceedings of the International Teletraffic Congress (ITC)*. Wurzburg, Germany.
- [64] Eben Howard, Clint Cooper, Mike Wittie, Steven Swinford, and Qing Yang. 2014. Cascading Impact of Lag on Quality of Experience in Cooperative Multiplayer Games. In *Proceedings of ACM Network and System Support for Games Workshop (NetGames)*. Nagoya, Japan.
- [65] Naoki Ishii, Yutaka Ishibashi, and Shinji Sugawara. 2009. Enhancement of Adaptive Delta-Causality Control with Adaptive Dead-reckoning for Multiplayer Online Games. In *Proceedings of ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (ACE)*. Athens, Greece.
- [66] Zenja Ivkovic, Ian Stavness, Carl Gutwin, and Steven Sutcliffe. 2015. Quantifying and Mitigating the Negative Effects of Local Latencies on Aiming in 3D Shooter Games. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. Seoul, Republic of Korea.
- [67] J. Clementa. 2021. Number of Video Gamers Worldwide in 2020. Online: <https://tinyurl.com/y23x256z>.
- [68] R. Jota J. Deber, C. Forlines, and D. Wigdor. 2015. How Much Faster is Fast Enough?: User Perception of Latency & Latency Improvements in Direct and Indirect Touch. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. Seoul, Republic of Korea.
- [69] Iryanto Jaya, Elvis S. Liu, and Youfu Chen. 2016. Combining Interest Management and Dead Reckoning: A Hybrid Approach for Efficient Data Distribution in Multiplayer Online Games. In *Proceedings of the 20th International Symposium on Distributed Simulation and Real-Time Applications (DS-RT)*. IEEE Press.
- [70] David R. Jefferson. 1985. Virtual Time. *ACM Transactions on Programming Language and Systems* 7, 3 (July 1985), 404–425. <https://doi.org/10.1145/3916.3988>
- [71] Xinbo Jiang, Farzad Safaei, and Paul Boustead. 2005. Latency and Scalability: a Survey of Issues and Techniques for Supporting Networked Games. In *Proceedings of the 13th IEEE International Conference on Networks (ICN)*. Kuala Lumpur, Malaysia.
- [72] Katie Jones. 2020. Online Gaming: The Rise of a Multi-Billion Dollar Industry. Visual Capitalist. Online: <https://tinyurl.com/y5l5ppmt>.
- [73] Arnaud Kaiser, Nadjib Achir, and Khaled Boussetta. 2010. Improving Energy Efficiency and Gameplay Fairness for Time-Sensitive Multiplayers Games in MANET. Cape Town, South Africa, 1 – 5. <https://doi.org/10.1109/ICCW.2010.5503904>
- [74] M. Katchabaw and R. Hanna. 2005. Bringing New HOPE to Networked Games: Using Optimistic Execution to Improve Quality of Service. In *Proceedings of the DiGRA Conference*. Vancouver, BC, Canada.
- [75] Vasily Kharitonov. 2012. Motion-Aware Adaptive Dead Reckoning Algorithm for Collaborative Virtual Environments. In *Proceedings of the 11th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry (VRCAI)*. Singapore.
- [76] Joohwan Kim, Pyarelal Knowles, Josef Spjut, Ben Boudaoud, and Morgan Mcguire. 2020. Post-Render Warp with Late Input Sampling Improves Aiming Under High Latency Conditions. In *Proceedings of the ACM High Performance Graphics Conference*. Virtual Conference.
- [77] Michael Kopietz. 2019. Systems And Methods for Player Input Motion Compensation by Anticipating Motion Vectors and/or Caching Repetitive Motion Vectors. US patent number 10,341,678 B2.
- [78] Yuji Kusunose, Yutaka Ishibashi, Norishige Fukushima, and Shinji Sugawara. 2012. Adaptive Delta-Causality Control with Prediction in Networked Real-Time Game Using Haptic Media. In *Proceedings of the 18th Asia-Pacific Conference on Communications (APCC)*. Jeju, Korea (South).
- [79] Anh Le and Yanni Ellen Liu. 2007. Fairness in Multi-Player Online Games on Deadline-Based Networks. In *Proceedings of the 4th IEEE Consumer Communications and Networking Conference (CCNC)*. IEEE Computer Society, Las Vegas, NV, USA, 670–675. <https://doi.org/10.1109/CCNC.2007.137>
- [80] Huy Viet Le, Valentin Schwind, Philipp Göttlich, and Niels Henze. 2017. PredicTouch: A System to Reduce Touchscreen Latency Using Neural Networks and Inertial Measurement Units. In *Proceedings of the ACM International Conference on Interactive Surfaces and Spaces (ISS)*. Brighton, UK.
- [81] LeagueBoards.net. [n.d.]. League of Legends Forum. <https://leagueboards.net/>.

- [82] Injung Lee, Sunjun Kim, and Byungjoo Lee. 2019. Geometrically Compensating Effect of End-to-End Latency in Moving-Target Selection Games. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Glasgow, Scotland, UK, 1–12.
- [83] Kyungmin Lee, David Chu, Eduardo Cuervo, Johannes Kopf, Yury Degtyarev, Sergey Grizan, Alec Wolman, and Jason Flinn. 2015. Outatime: Using Speculation to Enable Low-Latency Continuous Interaction for Mobile Cloud Gaming. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys)*. Florence, Italy, 151–165.
- [84] Steven Lee and Rocky Chang. 2017. On ‘Shot Around a Corner’ in First-person Shooter Games. In *Proceedings of the 15th Annual Workshop on Network and Systems Support for Games (NetGames)*. IEEE, Taipei, Taiwan, 1–6.
- [85] Steven W. K. Lee and Rocky K. C. Chang. 2018. Enhancing the Experience of Multiplayer Shooter Games via Advanced Lag Compensation. In *Proceedings of the ACM Multimedia Systems Conference (MMSys)*. Amsterdam, Netherlands, 284–293.
- [86] Wai-Kiu Lee and Rocky Chang. 2015. Evaluation of Lag-Related Configurations in First-Person Shooter Games. In *Proceedings of the Workshop on Network and Systems Support for Games (NetGames)*. IEEE Press, Zagreb, Croatia.
- [87] Shaolong Li and Changja Chen. 2006. Interest Scheme: A New Method for Path Prediction. In *Proceedings of 5th ACM SIGCOMM Workshop on Network and System Support for Games (NetGames)*. Association for Computing Machinery, Singapore.
- [88] Shaolong Li, Changjia Chen, and Lei Li. 2007. A Method of Using Personal Habits for Path Prediction in Network Games. In *The First International Symposium on Data, Privacy, and E-Commerce (ISDPE)*. Chengdu, China.
- [89] Shaolong Li, Changja Chen, and Lei Li. 2008. A New Method for Path Prediction in Network Games. *Computers in Entertainment* 5, 4, Article 8 (March 2008), 12 pages.
- [90] Zhi Li, Hugh Melvin, Rasa Bruzgiene, Peter Pocta, Lea Skorin-Kapov, and Andrej Zgank. 2018. Lag Compensation for First-Person Shooter Games in Cloud Gaming. In *Autonomous Control for a Reliable Internet of Services*. Springer International Publishing, 104–127.
- [91] Yow-Jian Lin, Katherine Guo, and Sanjoy Paul. 2002. Sync-MS: Synchronized Messaging Service for Real-time Multiplayer Distributed Games. Paris, France, 155–164. <https://doi.org/10.1109/ICNP.2002.1181396>
- [92] Shengmei Liu, Atsuo Kuwahara, James Scovell, Jamie Sherman, and Mark Claypool. 2021. L33t or N00b? How Player Skill Alters the Effects of Network Latency on First Person Shooter Game Players. In *Proceedings of the Game Systems Workshop (GameSys)*. Istanbul, Turkey.
- [93] Shengmei Liu, Atsuo Kuwahara, Jamie Sherman, James Scovell, and Mark Claypool. 2021. Lower is Better? The Effects of Local Latencies on Competitive First Person Shooter Game Players. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Virtual Conference.
- [94] M. Long and C. Gutwin. 2018. Characterizing and Modeling the Effects of Local Latency on Game Performance and Experience. In *Proceedings of the ACM Symposium on Computer-Human Interaction in Play (CHI Play)*. Melbourne, VC, Australia.
- [95] I. Scott MacKenzie and Colin Ware. 1993. Lag As a Determinant of Human Performance in Interactive Systems. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Amsterdam, Netherlands, 6 pages.
- [96] Regan L. Mandryk and Carl Gutwin. 2008. Perceptibility and Utility of Sticky Targets. In *Proceedings of Graphics Interface (GI)*. Windsor, Ontario, Canada.
- [97] Martin Mauve. 2000. Consistency in Replicated Continuous Interactive Media. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*. Philadelphia, PA, USA.
- [98] Martin Mauve. 2000. How to Keep a Dead Man from Shooting. In *Proceedings of Interactive Distributed Multimedia Systems and Telecommunication Services Workshop (IDMS)*. Enschede, Netherlands.
- [99] Martin Mauve, Jürgen Vogel, Volker Hilt, and Wolfgang Effelsberg. 2004. Local-lag and Timewarp: Providing Consistency for Replicated Continuous Applications. *IEEE Transactions on Multimedia* 6, 1 (2004), 47–57.
- [100] W. McCarty, S. Sheasby, C. Switzer, P. Amburn, and M. R. Stytz. 1994. A Virtual Cockpit for a Distributed Interactive Simulation. *IEEE Computer Graphics and Applications* 14, 01 (jan 1994).
- [101] Teemu Mäki-patola and Perttu Hämäläinen. 2004. Latency Tolerance for Gesture Controlled Continuous Sound Instrument Without Tactile Feedback. In *Proceedings of the International Computer Music Conference (ICMC)*. Coral Calbles, FL, USA.
- [102] A. Ng, M. Annett, P. Dietz, A. Gupta, and W. Bischof. 2014. In the Blink of an Eye: Investigating Latency Perception During Stylus Interaction. In *In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. Toronto, Canada.
- [103] Doowon Paik, Chung-Ha Yun, and Jooyeon Hwang. 2008. Effective Message Synchronization Methods for Multiplayer Online Games with Maps. *Computuer Humman Behavior* 24, 6 (Sept. 2008), 9 pages.

- [104] Wladimir Palant, Carsten Griwodz, and Pål Halvorsen. 2006. Evaluating Dead Reckoning Variations with a Multi-Player Game Simulator. In *Proceedings of the International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*. Association for Computing Machinery, Newport, RI, USA, Article 4, 6 pages.
- [105] Lothar Pantel and Lars C. Wolf. 2002. On the Impact of Delay on Real-Time Multiplayer Games. In *Proceedings of the Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*. Miami, FL, USA.
- [106] Lothar Pantel and Lars C. Wolf. 2002. On the Suitability of Dead Reckoning Schemes for Games. In *Proceedings of the 1st Workshop on Network and System Support for Games (NetGames)*. Braunschweig, Germany.
- [107] A. Pavlovych and C. Gutwin. 2012. Assessing Target Acquisition and Tracking Performance for Complex Moving Targets in the Presence of Latency and Jitter. In *Proceedings of Graphics Interface*. Toronto, ON, Canada.
- [108] Ricky Pusch. 2019. Explaining How Fighting Games Use Delay-Based and Rollback Netcode. arsTECHNICA. <https://tinyurl.com/y69jt5fz>.
- [109] Peter Quax, Patrick Monsieurs, Wim Lamotte, Danny De Vleschauwer, and Natalie Degrande. 2004. Objective and Subjective Evaluation of the Influence of Small Amounts of Delay and Jitter on a Recent First Person Shooter Game. In *Proceedings of ACM Network and System Support for Games Workshop (NetGames)*. Portland, OR, USA.
- [110] Kjetil Raaen and Andreas Petlund. 2015. How Much Delay Is There Really in Current Games? *Proceedings of ACM Multimedia Systems (MMSys)*.
- [111] Jared Ramsey. 2020. Online Gaming: The Rise of a Multi-Billion Dollar Industry. Lineups. Online: <https://www.lineups.com/esports/top-10-esports-games/>.
- [112] Brent Randall. 2020. Valorant's 128-Tick Servers. Riot Games. <https://technology.riotgames.com/news/valorants-128-tick-servers>.
- [113] David Roberts, Damien Marshall, Séamus Mcloone, Declan Delaney, Rob Aspin, and Tomas Ward. 2008. Bounding Inconsistency Using a Novel Threshold Metric for Dead Reckoning Update Packet Generation. *Simulation* 84, 5 (05 2008), 239–256. <https://doi.org/10.1177/0037549708092221>
- [114] Saeed Shafiee Sabet, Steven Schmidt, Saman Zadtootaghaj, Carsten Griwodz, and Sebastian Moller. 2018. Towards Applying Game Adaptation to Decrease the Impact of Delay on Quality of Experience. In *IEEE International Symposium on Multimedia (ISM)*. Taichung, Taiwan, 114–121.
- [115] Saeed Shafiee Sabet, Steven Schmidt, Saman Zadtootaghaj, Babak Naderi, Carsten Griwodz, and Sebastian Möller. 2020. A Latency Compensation Technique Based on Game Characteristics to Mitigate the Influence of Delay on Cloud Gaming Quality of Experience. In *Proceedings of the ACM Multimedia Systems Conference (MMSys)*. Istanbul, Turkey.
- [116] Robert Salay and Mark Claypool. 2020. A Comparison of Automatic versus Manual World Alteration for Network Game Latency Compensation. In *Proceedings of the ACM Annual Symposium on Computer-Human Interaction in Play (CHI PLAY)*. Virtual Conference. Work In Progress.
- [117] Cheryl Savery, Nicholas Graham, Carl Gutwin, and Michelle Brown. 2014. The Effects of Consistency Maintenance Methods on Player Experience and Performance in Networked Games. In *Proceedings of the ACM Conference on Computer Supported Cooperative work (CSCW)*. Baltimore, MD, USA.
- [118] Cheryl Savery, Nicholas Graham, Carl Gutwin, and Michelle Brown. 2014. The Effects of Consistency Maintenance Methods on Player Experience and Performance in Networked Games. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. Baltimore, MD, USA, 1344–1355.
- [119] Cheryl Savery and T. C. Graham. 2013. Timelines: Simplifying the Programming of Lag Compensation for the next Generation of Networked Games. *Multimedia Systems* 19, 3 (June 2013), 271–287. <https://doi.org/10.1007/s00530-012-0271-3>
- [120] Cheryl Savery, T. C. Nicholas Graham, and Carl Gutwin. [n.d.]. The Human Factors of Consistency Maintenance in Multiplayer Computer Games. In *Proceedings of the 16th ACM International Conference on Supporting Group Work*. Association for Computing Machinery, Sanibel Island, FL, USA, 187–196.
- [121] Jörg R.J. Schirra. 2001. Content-Based Reckoning for Internet Games. In *GAME-ON - the Simulation and AI in Games Conference*. London, UK.
- [122] P. M. Sharkey, M. D. Ryan, and D. J. Roberts. 1998. A Local Perception Filter for Distributed Virtual Environments. In *Proceedings of the Virtual Reality Annual International Symposium (VRAIS)*. Atlanta, GA, USA.
- [123] Nathan Sheldon, Eric Girard, Seth Borg, Mark Claypool, and Emmanuel Agu. 2003. The Effect of Latency on User Performance in Warcraft III. In *Proceedings of ACM Network and System Support for Games Workshop (NetGames)*. Redwood City, CA, USA.
- [124] Gabriel Shelley and Michael Katchabaw. 2005. Patterns of Optimism for Reducing the Effects of Latency in Networked Multiplayer Games. In *Proceedings of FuturePlay*. Lansing, MI, USA.

- [125] H. Shen and Suiping Zhou. 2013. Achieving critical consistency through progressive slowdown in highly interactive Multi-Player Online Games. In *Proceedings of the IEEE 17th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*. Whistler, BC, Canada.
- [126] J. Smed, T. Kaukoranta, and H. Hakonen. 2002. Aspects of Networking in Multiplayer Computer Games. *The Electronic Library* 20, 2 (2002), 87–97.
- [127] Jouni Smed, Timo Kaukoranta, and Harri Hakonen. 2002. Aspects of Networking in Multiplayer Computer Games. *The Electronic Library* 20, 2 (2002), 87–97.
- [128] Jouni Smed, Henrik Niinisalo, and Harri Hakonen. 2004. Realizing Bullet Time Effect in Multiplayer Games with Local Perception Filters. In *Proceedings of 3rd ACM Workshop on Network and System Support for Games (NetGames)*. Portland, OR, USA, 8 pages.
- [129] Richard H Y So and Micheal J. Griffin. 1992. Compensating Lags in Head-coupled Displays Using Head Position Prediction and Image Deflection. *Journal of Aircraft* 29, 6 (January 1992), 1064–1068. <https://eprints.soton.ac.uk/429168/>
- [130] Josef Spjut, Ben Boudaoud, Kamran Binaee, Jonghyun Kim, Alexander Majercik, Morgan McGuire, David Luebke, and Joohwan Kim. 2019. Latency of 30 ms Benefits First Person Targeting Tasks More than Refresh Rate above 60 Hz. In *Proceedings of SIGGRAPH Asia*. Brisbane, QLD, Australia, 110–113.
- [131] Michael Stallone. 2019. 8 Frames in 16ms: Rollback Networking in Mortal Kombat and Injustice 2. YouTube. https://www.youtube.com/watch?v=7jb0FOcImdg&feature=youtu.be&ab_channel=GDC.
- [132] steamcommunity.com. [n.d.]. rl Forum. <https://steamcommunity.com/app/252950/discussions/>.
- [133] David Straily and Dave Heironymus. 2020. Netcode and 128-Servers. YouTube. https://www.youtube.com/watch?v=_Cu97mr7zcM.
- [134] Dane Stuckel and Carl Gutwin. 2008. The Effects of Local Lag on Tightly-Coupled Interaction in Distributed Groupware. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)*. San Diego, CA, USA.
- [135] Jiawei Sun and Mark Claypool. 2019. Evaluating Streaming and Latency Compensation in a Cloud-based Game. In *Proceedings of the 15th IARIA Advanced International Conference on Telecommunications (AICT)*. Nice, France.
- [136] Neema Teymory. 2016. Why Making Multiplaye Games is Hard: Lag Compensating Weapons in MechWarrior Online. Gamasutra. <https://tinyurl.com/gamasutra-teymory-2016>.
- [137] Alexey Tumanov, Robert Allison, and Wolfgang Stuerzlinger. 2007. Variability-Aware Latency Amelioration in Distributed Environments. In *IEEE Virtual Reality Conference*. Charlotte, NC, USA, 123–130. <https://doi.org/10.1109/VR.2007.352472>
- [138] Rosane Ushirobira, Denis Efimov, Géry Casiez, Nicolas Roussel, and Wilfrid Perruquetti. 2016. A Forecasting Algorithm for Latency Compensation in Indirect Human-computer Interactions. In *European Control Conference (ECC)*. IEEE, Aalborg, Denmark, 1081–1086.
- [139] Ivan Vaghi, Chris Greenhalgh, and Steve Benford. 1999. Coping with Inconsistency Due to Network Delays in Collaborative Virtual Environments. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*. London, UK.
- [140] Rodrigo Vicencio-Moreira, Regan L. Mandryk, Carl Gutwin, and Scott Bateman. 2014. The Effectiveness (or Lack Thereof) of Aim-Assist Techniques in First-Person Shooter Games. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Toronto, ON, Canada.
- [141] Greger Wikstrand, Lennart Schedin, and Fredrik Elg. 2004. High and Low Ping and the Game of Pong Effects of Delay and Feedback. In *Proceedings of 3rd ACM Workshop on Network and System Support for Games (NetGames)*. Portland, OR, USA.
- [142] Jiann-Rong Wu and Ming Ouhyoung. 2000. On Latency Compensation and Its Effects for Head Motion Trajectories in Virtual Environments. *The Visual Computer* 16 (03 2000), 79–90.
- [143] Jingxi Xu and B. Wah. 2013. Concealing Network Delays in Delay-sensitive Online Interactive Games based on Just-noticeable Differences. In *IEEE International Conference on Multimedia and Expo (ICME)*. San Jose, CA, USA, 1–6.
- [144] Amir Yahyavi, Kévin Huguenin, and Bettina Kemme. 2012. Interest Modeling in Games: The Case of Dead Reckoning. *Multimedia Systems* 19 (06 2012). <https://doi.org/10.1007/s00530-012-0275-z>
- [145] Seok-Jong Yu and Yoon-Chul Choy. 2001. A Dynamic Message Filtering Technique for 3D Cyberspaces. *IEEE Computer Communications* 24 (2001).
- [146] Sebastian Zander, Ian Leeder, and Grenville Armitage. 2005. Achieving Fairness in Multiplayer Network Games through Automated Latency Balancing. In *Proceedings of the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (ACE)*. Valencia, Spain.
- [147] Mao-Jun Zhang and Nicolas Georganas. 2004. An Orientation Update Message Filtering Algorithm in Collaborative Virtual Environments. *Journal of Computer Science and Technology* 19 (2004).

- [148] Yi Zhang, Ling Chen, and Gencai Chen. 2006. Globally Synchronized Dead-Reckoning with Local Lag for Continuous Distributed Multiplayer Games. In *Proceedings of the ACM Workshop on Network and System Support for Games (NetGames)*. Singapore.
- [149] Sili Zhao, Du Li, Hansu Gu, Bin Shao, and Ning Gu. 2009. An Approach to Sharing Legacy TV/Arcade Games for Real-Time Collaboration. In *Proceedings of the 29th IEEE International Conference on Distributed Computing Systems (ICDCS)*. IEEE Computer Society, Montreal, QC, Canada, 165–172.