

On Consistency and Network Latency in Distributed Interactive Applications: A Survey – Part I

Authors:

Declan Delaney

Tomas Ward

Seamus McLoone

Affiliations:

Department of Electronic Engineering

National University of Ireland, Maynooth

Maynooth

Co. Kildare

Ireland

Corresponding Author details:

Declan Delaney

Via Matteotti 9,

20097 San Donato Milanese,

Italy

Email: drd@eircom.net

Tel: 0039 02 55600325 or 0039 02 52061478

On Consistency and Network Latency in Distributed Interactive Applications: A Survey – Part I

Declan Delaney, Tomas Ward, Seamus McLoone

Abstract - This paper is the first part of a two-part paper that documents a detailed survey of the research carried out on consistency and latency in distributed interactive applications (DIAs) in recent decades. Part I reviews the terminology associated with DIAs and offers definitions for consistency and latency. Related issues such as jitter and fidelity are also discussed. Furthermore, the various consistency maintenance mechanisms that researchers have used to improve consistency and reduce latency effects are considered. These mechanisms are grouped into one of three categories, namely time management, Information management and system architectural management. This paper presents the techniques associated with the time management category. Examples of such mechanisms include time warp, lock step synchronisation and predictive time management. The remaining two categories are presented in part two of the survey.

I. INTRODUCTION

Human interaction with computers has evolved from centralized batch processing using Hollerith cards to virtual reality systems in which users can be fully immersed. The origins of virtual reality technology can be traced back to vehicle simulation research in the 1920s (Ellis 1994) and since then, virtual environments have spread to areas such as military simulations, cooperative whiteboards and architectural design.

The first distributed software virtual environment was SIMNET, a United States research program initiated in 1983 to train soldiers in battlefield tactics (Calvin, Dickens, Gaines, Metzger, Miller et al. 1993) and which culminated in the development of the DIS protocol standard in 1993 (IEEE 1993). Since then numerous academic, military and commercial distributed applications have been developed and documented (Joslin, Di Giacomo and magnenat-Thalman 2004). The diverse applications that exploit distributed virtual technology have been referred to using various terms. In this paper they will be referred to as Distributed Interactive Applications or DIAs.

DIAs are subject to many problems. However, the authors believe that the key common objective running through most of the research in the area of DIAs is related to the issue of consistency and the maintenance of adequate consistency among all participants of the DIA in spite of demands on system resources. In addition they believe that the single greatest contributing factor to spatial and temporal inconsistencies experienced by end users in the virtual world is network latency or lag.

The issue of consistency has been the subject of much research in the field of distributed systems for the past few decades. However, no single detailed survey of these techniques has been produced, although some publications summarize a number of the issues and approaches taken: (Snowdon, Greenhalgh, Benford, Bullock and Brown 1996; Macedonia and Zyda 1997; Singhal and Zyda 1999; Joslin et al. 2004; Roberts 2004). This paper aims to bridge this gap and provide a concise summary of the issues and the approaches taken to tackle the problem of maintaining consistency in the presence of network latency. It also reviews the terminology associated with the domain of distributed interactive applications and proposes definitions for the terms consistency and latency. This survey is not intended to be comprehensive in bibliographic terms; instead references are selected to best illustrate, in our opinion, the different concepts.

The next section looks at some basic terminology. This is provided for ease of reference. Section III investigates what is meant by the term *distributed interactive application*. Sections IV and V define and review the inherent problems of consistency and network latency in such applications. In section VI we provide a classification for various mechanisms that have been employed to mask network latency and improve consistency in DIAs. The first of these classifications, time management techniques, is then presented. The remaining two classifications are presented in part II of the survey.

II. TERMINOLOGY

To facilitate discussion of DIAs, the following terms derived from the world of modeling and simulation are listed (Zeigler, Praehofer and Kim 2000). The terms listed here are used extensively in referring to DIAs, but are rarely defined. Other terms will be defined as they are encountered.

State: a complete description of a virtual entity at a single moment in time (Churchill, Snowdon and Munro 2001). In a DIA information pertaining to state variations in the virtual environment should be shared among all participants. This information is often referred to as Dynamic Shared State (Singhal et al. 1999; Qin 2002).

Entity or Object: an element of the synthetic environment that is created and controlled by a simulation application through the exchange of information (IEEE 1993). We consider entity and object as synonymous, and we adopt the term entity throughout this paper. Entities may be *passive* or *active* (Singhal et al. 1999). A passive entity is one that is either entirely stationary or moves deterministically. The behavior of an active entity is non-deterministic, although it may be predictable in the sense that its motion is a function of the constraints imposed by the environment, the objectives/abilities of the participant and the initial conditions. A similar

classification describes entities as either *static* (passive) or *dynamic* (active) (Lui 2001). The state of a static entity never changes with time whereas the state of a dynamic entity can.

Environment: the information needed to render an application's time-constant state (Mauve, Hilt, Kuhmunch and Effelsberg 2001).

Event: The state of a DIA may change for two reasons, as a result of the *passage of time* or the occurrence of *events*. An event causes a state change that is not a fully deterministic function of time (Mauve 2000). Between successive events the state of the medium is a fully deterministic function of time. Events can be separated into *external events* and *internal events*. External events are caused by (user) interactions with the medium, whereas Internal events are non-deterministic internal changes in the state of the application. Events may also be viewed as *deterministic* or *non-deterministic* (Roberts and Sharkey 1997; Sharkey, Ryan and Roberts 1998). Deterministic events can be fully predicted whereas non-deterministic events cannot be reliably predicted because they result from real-time interaction between the DIA and human participants.

Node: A node refers to a computing device connected to the communications network. A node can be a source or destination machine, a router or any other device that processes data.

Dead Reckoning: Dead reckoning exploits information about current user dynamics to make a short-term prediction of user movement based on extrapolation techniques. It is formally defined in the DIS protocol (IEEE 1993). It was first implemented as the Players and Ghosts paradigm in the VERN (Virtual Environment Realtime Network) test bed, developed by Blau et al. (1992).

III. DISTRIBUTED INTERACTIVE APPLICATION

A Distributed Interactive Application (DIA) is best defined by considering the three terms that comprise its title:

1. A system is *distributed* if there is a message transmission delay between parts of the system that is not negligible compared to the time between events in a single process (Lamport 1978);
2. *Interactive* refers to an input-output process involving input from a user through a human-computer-interface with an appropriate output response from the system (Broll 1997; Natrajan and Reynolds Jr 1999; Manninen 2000; Zhou, Cai, Lee and Turner 2001). Interaction is also defined as truly concurrent object manipulation (Margery, Arnaldi and Plouzeau 1999), as distinct from the cooperative sequential object manipulation defined in (Broll 1997).
3. The *application* is a software system built for a specific purpose. It is usually a virtual environment that, in some way, reflects the physical world and interfaces to the physical world.

Distributed Interactive Simulation (DIS)
Networked Virtual Environment (NetVE) (Singhal et al. 1999)
Distributed Virtual Environment (DVE) (Stytz 1996)
Distributed Interactive Media (DIM) (Mauve et al. 2001)
Networked Interactive Entertainment (NIE) (Capps, McDowell and Zyda 2001)
Collaborative Virtual Environment (CVE) (Park and Kenyon 1999; Vaghi, Greenhalgh and Benford 1999)
Distributed Synthetic Environments (DSE) (Worthington and Roberts 2000)
Shared Virtual Environment (SVE) (Waters and Barrus 1997)
Computer Supported Cooperative Work (CSCW) (Greif 1988)
Groupware Systems (Ellis and Gibbs 1989)

Table 1: DIAs are referred to using various terms and acronyms

What this paper refers to as DIAs have been referred to using many different terms, as summarized in Table 1. These reflect the diverse applications of this paradigm and the different research emphasis of each research group: shared workspaces, networked games, distributed whiteboards, group editors, distributed architectural design, education, telemedicine and simulations (Bhola, Banavar and Ahamad 1998; Sun, Jia, Zhang, Yang and Chen 1998; Bouras, Hornig, Triantafillou and Tsiatsos 2001; Fujimoto 2001; Riva and Gamberini 2001; Tawfik and Fernando 2001; McCoy, Delaney and Ward 2003; Frecon 2004). Several research teams have developed experimental distributed virtual environment platforms; some examples include RING (Funkhouser 1995), NPSNET (Macedonia, Zyda, Pratt, Barham and Zeswitz 1994; Capps, McGregor, Brutzman and Zyda 2000), MASSIVE (Greenhalgh and Benford 1995; Greenhalgh, Purbick and Snowdon 2000), PaRADE (Roberts, Sharkey and Sandoz 1995; Roberts and Sharkey 1997; Roberts et al. 1997), SPLINE (Barrus, Waters and Anderson 1996), CAVERNSoft (Leigh, Yu, Schonfeld, Ansari, He et al. 2001), VELVET (de Oliveira and Georganas 2002), PARADISE (Holbrook, Singhal and Cheriton 1995; Singhal 1996), DIVE (Frécon and Stenius 1998; Frécon 2003; Frécon 2004), QUICK (Capps 2000), VPark (Joslin, Molet, Thalmann, Esmerado, Thalmann et al. 2001), MOVE (Garcia, Montala, Pairot, Rallo and Skarmeta 2002), ATLAS (Lee, Lim and Han 2002), EQUATOR (MacColl, Millard, C. and Steed 2002) and PING (Roberts 2004).

For the purposes of this paper we adopt the term Distributed Interactive Application (Kelly 1997; Diot and Gautier 1999; Lee, Yang, Yoon, Yu and Hyun 2000; Zhou et al. 2001; Qin 2002; Vogel, Mauve, Hilt and Effelsberg. 2003) as we consider it to be a generic term that encapsulates all aspects of networked applications that involve interaction, collaboration, cooperation and communication . We define a DIA to be:

a networked software system that seeks to maintain global consistency when responding to multiple simultaneous non-deterministic inputs. (1)

A key aspect of a DIA is the participation of humans-in-the-loop (Ellis et al. 1989). People interact with the system in real-time and they expect the system to respond in real-time to the actions they make. In addition they assume that other users of the system, particularly those with whom they are interacting, share the same view of the system state. In other words, users demand consistency. The next section addresses this issue in detail.

IV. CONSISTENCY

The input from users of a DIA generates events that then modify an underlying database that is shared in some way by all users across a communications network. An essential requirement of the DIA is that users are informed of changes to this common database in real time if the changes affect them or if they are interested in receiving the changes. Users of DIAs expect both temporal and spatial consistency so that their experience of the virtual world can be the same as the users with whom they are interacting.

Consistency is one of the key factors in providing a truly interactive experience to users of DIAs (Vaghi et al. 1999; Bowman, Johnson and Hodges 2001; Gutwin 2001; Henderson 2003). So, what exactly is consistency? For Gautier et al. (1999) it implies that at any point in time, all players should ideally see the same information at the same time independent of the network. This is referred to as absolute consistency and it has been defined mathematically by Qin (2002) and Zhou (2001). Absolute consistency in distributed applications is impossible to attain because of the existence of network latency and the fact that the delay in transmitting information across the application cannot be ignored when compared to the time between events.

Consistency refers to a number of aspects in DIAs (Shneiderman 1984; Dourish 1995; Blakowski and Steinmetz 1996; Roberts 2004):

1. Synchronisation: the maintenance of (a) temporal relations between events so that the time of each event relative to other events across the DIA is the same for all participants and (b) spatial relations so that entity positions are the same across the DIA;
2. Causality or Ordering: events cause the state of the system to change to a new state. The order in which events are received is thus important to maintain a natural cause-effect order;
3. Concurrency: the simultaneous execution of events by different users on the same entity within the application. Entity ownership conflicts have to be resolved.

Linked to all of these aspects are the issues of responsiveness and fidelity. Responsiveness is the time taken for the system to register and respond to a user event. Fidelity is defined in the DIS standard as the degree to which the representation within a simulation is similar to a real-world object, feature or condition in a measurable or perceivable manner (IEEE 1995). Similarly the Fidelity Definition and Metrics Implementation Study Group (DM-ISG) defines fidelity as the degree to which a model or simulation reproduces the state and behaviour of a real-world object or perception of a real-world object, feature, condition or standard in a measurable or perceivable manner.” The DM-ISG identifies six components to fidelity: resolution, error, accuracy, sensitivity, precision and capacity (Roza, Voogd, Jense and van Gool 1999). Stytz (1996) explores fidelity in DIAs and lists a number of types of fidelity, such as physics fidelity, time fidelity and sensory fidelity. Capps (2000) discusses the link between visual accuracy and

fidelity and believes that visual accuracy is only a 'passing first-order approximation for fidelity.' Fidelity is also referred to as *level-of-detail* or *resolution* (Radhakrishnan and Wilsey 1999).

Poor consistency can lead to fidelity problems and fidelity is often sacrificed to maintain consistency and responsiveness. In an ideal DIA absolute consistency could be obtained: all user clocks would be synchronized, all events would be executed in the order they occurred, conflicts over shared entities and the system response time would satisfy normal user expectations, and fidelity would be guaranteed. . However, this ideal cannot be achieved and the inconsistencies that result manifest themselves in many different ways, resulting in three noted phenomena (Schwarz and Mattern 1994; Dourish 1995; Sun et al. 1998; Vaghi et al. 1999; Zhou et al. 2001; Zhou, Cai, Lee and Turner 2003):

1. **Divergence:** this refers to the final temporal-spatial state of the environment being different for different participants. Users will thus be unwittingly be interacting with an inconsistent environment so that their behaviour will also be inconsistent, possibly causing the states to diverge even further;
2. **Causality violation:** events may be received, executed and rendered out of their natural cause-effect order, so that users notice the effect before the cause;
3. **Intention / expectation violation:** the actual effect of an action may differ from the intended effect because events may be generated concurrently. This refers to something that is expected to happen based on real-world experience, but doesn't happen in the virtual world.

Consistency therefore needs to be controlled within the DIA. This is achieved by using consistency maintenance mechanisms. In this paper these are defined as:

any element employed to ensure a sufficient, uniform dynamic shared state for all participants in a DIA. (2)

While absolute consistency is impossible to achieve in practical DIAs because of their dynamic nature, a hypothetical DIA in which the environment remains totally static and unchanging would achieve absolute consistency. At one extreme therefore is the dynamic DIA with continuous state changes that requires a high throughput of update packets and that never achieves absolute consistency; at the other is the totally static DIA where no update packets are required and so absolute consistency can be achieved. This is known as the *Consistency-Throughput Trade-off* (Singhal et al. 1999). There is also a conflict between system responsiveness and consistency. High responsiveness can be achieved locally, but with the risk that the global state of the DIA becomes inconsistent because of network latency (Cheshire 1996). This conflict has been documented as the consistency-responsiveness trade-off (Zhou et al. 2001; Mauve, Vogel, Hilt and Effelsberg 2002; Qin 2002).

The principal cause of inconsistencies is ironically linked to the network that actually facilitates DIAs. Therefore the relationship that exists between the physical network and consistency needs to be explored. In particular we focus on the most important aspect of the physical network affecting consistency: latency. The next section describes this phenomenon.

V. LATENCY

A review of the research literature reveals that a commonly agreed definition of latency does not exist. Latency has sometimes been defined as the amount of time required to transfer a bit of data from one point to another (Singhal et al. 1999). This definition leaves the word *point* open to several possible interpretations. Latency is also defined in (Smed, Kaukoranta and Hakonen

2001) as the length of time (or delay) that a message incurs during transmission from one designated node to another. Again, the definition is open to different interpretations of the word *node*. A similar problem exists in the following definition: the time taken by data to travel from the source to the destination (Dutta-Roy 2000). Meehan et al. incorporate the delay associated with graphics rendering in their definition (Meehan, Razzaque, Whitton and Brooks 2003). Tse-Au and Morreale interpret latency as the time required for one bit to propagate round-trip between two nodes, similar to a *ping* time (Tse-Au and Morreale 2000). Pullen and Wood define it as the time delay that occurs between the output of a data packet at the application level of one simulator and the input of that data packet at the application level of another simulator (Pullen and Wood 1995). We have taken this definition as a basis to propose the following definition for network latency based on the ISO OSI reference model:

Network Latency is the time taken from the start of exchange of an application protocol data unit (APDU) at the application layer of one participating node to the end of the exchange of the same APDU with the application layer of a second participating node.

(3)

This definition is similar to that put forward by de Oliveira et al. (1999) and excludes latency due to any other source, for example, computations in the application, graphics rendering or user reaction time. These other sources are also included in local latency, which is the latency involved in representing a user's actions back to that user (Roberts et al. 1995). For the rest of the paper the term latency will refer to network latency.

Latency is particularly problematic in interactive applications where the network delays are comparable to the interaction time or speed (Sharkey et al. 1998). Typical latency values to

maintain real-time interaction fluctuate between 40 and 300ms (Diot et al. 1999). Cheshire (1996) suggests that the round-trip latency of 100ms should be the maximum delay, whereas both Wloka (1995) and Diot and Gautier (1999) aim for a one-way delay of 100ms. Conscious awareness of latency by users of a DIA depends on the application. For voice interaction a maximum latency of 100ms is sufficient (Cheshire 1996) and Mauve et al. report values of 120ms for visual data (Mauve et al. 2002). In teleoperation control loops, instability occurs when total system latency exceeds 100ms. For motor-driven tasks human reaction time is in the order of 200ms (Krumm-Heller and Taylor 2000).

Components of Latency

To assist in the understanding of latency, it is convenient to analyse the factors that comprise it. The main contributors to latency are (Martin, McGregor and Cleary 2000):

1. Queuing and processing at routers, bridges and gateways within the network and at source and destination nodes. Packet loss and timeouts may be considered to introduce infinite latency;
2. Transcoding delays associated with encryption / decryption of information and compression / decompression of data;
3. Propagation and transmission delays due to the finite speed of light and the speed of the communications link.

This paper decomposes latency into three types according to the source of the delay:

1. Packet processing delay;
2. Bit propagation delay;

3. Packet propagation delay.

Figure [1] here

The latency τ_{total} associated with a single packet can therefore be expressed as the sum of these three components. This is illustrated in Figure 1 and is represented by the following equation:

$$\tau_{total} = \sum_{i=1}^N \tau_{process}^i + \sum_{i=1}^{N-1} \Delta\tau_i^{i+1} + \sum_{i=1}^{N-1} \frac{M}{B_i^{i+1}} \quad (4)$$

\uparrow
*Packet
Processing
Delay*

\uparrow
*Bit
Propagation
Delay*

\uparrow
*Packet
Propagation
Delay*

where

$\tau_{process}^i$ is the time to process a packet at node i ;

N is the number of nodes, including source and destination nodes;

$\Delta\tau_i^{i+1}$ is the transmission time between nodes i and $i+1$;

B_i^{i+1} is the bandwidth between nodes i and $i+1$;

M is the number of bits in the packet.

The *packet processing delay* refers to the time taken to manage and process the data as it migrates through the network hardware and to process and parse the data at both source and destination nodes. This includes compressing, decompressing, encrypting, decrypting and any processing performed by the operating system or network hardware at the end-point computers. It also includes the time delay associated with flow control and congestion control, buffering and packet queuing.

The *bit propagation delay* refers to the delay associated with the physical speed of transmission (about $5\mu\text{s}$ per km).

The *packet propagation delay* refers to the time required for all bits in a packet to be transmitted across the network from source to destination node considering only the inter-node bandwidth.

From this division of latency the following may be noted:

1. The bit propagation delay cannot be eliminated and it has a lower theoretical limit imposed by the speed of light in a vacuum;
2. The packet processing delay can be reduced in a number of ways: by reducing the quantity of data on the network, by increasing the processing power at routers and source/destination nodes and by using more efficient processing algorithms;
3. The packet propagation delay can be reduced by increasing the available network bandwidth and by reducing the amount of data that must be transmitted between nodes.

Points (2) and (3) explain why most of the techniques and approaches that have been explored to combat the effects of network latency focus on reducing the quantity of network traffic and

improving software architectures to reduce both the packet processing delay and the packet propagation delay.

The division of latency given here facilitates a more accurate understanding of latency and the impact of latency reducing techniques and is similar to that espoused by Jehaes et al. (2003). Other divisions of latency have been presented elsewhere. Tse-Au and Morreale (2000) view latency as being tunable or un-tunable. *Un-tunable latency* refers to the sum of the signal propagation delay and the equipment processing delay in transmitting one bit across the network. *Tunable latency* is the time taken for one bit of data to propagate through all the queuing delays between the two communicating nodes. In (Hecker and Simpson 2001) latency is split between *path latency* and *queuing latency*. Path latency is the time taken for a message to get from one place to another. Queuing latency is the resultant delay when the data sent exceeds the bandwidth. Two kinds of latency are also described by Mauve et al. (2004): *observation lag* and *influence lag*. Observation lag is the delay between an event's occurrence and its display, whereas influence lag is the delay between our attempt to influence the world and the time taken for the influence to actually occur.

Network latency varies with prevailing network conditions. This variation in latency results in a phenomenon called jitter, which is described in the following paragraph.

Jitter

All information transmitted across the network is carried in the form of data packets. Each packet may follow a different route through the network and the order of arrival of packets at the destination node is not guaranteed by the network layer. The network is heterogeneous in nature with continuous variations in router queue lengths, buffering times, and routing paths for packets. Packet sizes also vary greatly and similarly the time to process packets as they migrate through

the network and reorder packets to rebuild complete packets at destination also vary significantly (Dutta-Roy 2000). Depending on network conditions, this may result in a different latency for each transmitted packet and gives rise to a phenomenon known as jitter (Blow 1998; Smed et al. 2001). For the purposes of this paper the following definition will be used:

Jitter refers to the unpredictable variation in latency with time. (5)

The effects of network latency and jitter on end-users were examined in an interesting suite of qualitative experiments performed by Vaghi et al.(1999). They examined the effects of gradually increasing network delays for one of the two participants in a networked virtual ball game to establish how delays affect the performance of the players, how performance breakdowns manifest themselves and how players adopt to delays. Park and Kenyon (1999) assessed the effects of latency and jitter on executing a cooperative task and found that from a psycho-perceptual viewpoint, jitter has the greatest impact when latency is high and the collaborative task is difficult. Jitter has a much greater impact than latency on performance. Gutwin (2001) investigated the effects of latency and jitter on two types of user interaction: prediction of movement and moving a shared object. He found that both performance and user strategy were affected. Other work has been carried out by Henderson (2003), who investigated how users adapt to latency in distributed games, Bowman et al. (2001), who analysed interaction in virtual environments and Mauve (2000), who examined the effects on interaction in a 3D tele-cooperation application when local lag is introduced to offset the effects of jitter.

Latency prevents the possibility of achieving an absolute consistent dynamic shared state across a DIA. As the quantity of data to be managed by the DIA increases (an issue called scalability (Macedonia 1995) and due, for example, to an increase in the number of active entities

and participants or the spatial extent of the DIA or the complexity of the tasks involved), maintaining consistency across the DIA becomes even more difficult. The following section examines the various methods, mechanisms and approaches for controlling consistency in DIAs in the face of this inherent network latency.

VI. CONSISTENCY MAINTENANCE MECHANISMS

A DIA can be viewed as a distributed database with multiple simultaneous users modifying it in real-time. The key problem then becomes that of ensuring that the database is consistent for all users, so that they are interacting with an up-to-date source of information. Consistency maintenance mechanisms refer to any element employed to ensure a sufficient, uniform dynamic shared state for all participants in a DIA. Consistency maintenance mechanisms are also referred to as latency hiding techniques (Cheshire 1996), consistency control algorithms (Roberts et al. 1997), distributed synchronization mechanisms (Diot et al. 1999) and object synchronization (Lui 2001). Of the many existing mechanisms, most of them are bandwidth-saving mechanisms and combat latency by reducing the number of packets being sent across the network as explained previously. Each consistency maintenance mechanism falls into one of two general categories (Jefferson 1990; Blanchard and Lake 1997; Fujimoto 2001; Cronin, Filstrup, Kurc and Jamin 2002):

1. **Optimistic/Aggressive:** this is a mechanism that takes risks by performing speculative computation, which, if subsequently determined to be correct, saves time, but if incorrect must be rolled back and corrected;

2. **Pessimistic/Conservative:** this is a mechanism that never indulges in speculative computation and hence never has to roll back. These algorithms perform poorly in fast-paced interactive applications where a constant rate of simulation is important.

Systems can employ a hybrid approach, applying optimistic and conservative mechanisms at different points in the DIA, depending on the type of DIA. The various mechanisms may be further classified into three general classes:

1. **Time Management Techniques** – These all manipulate time to mask the effects of latency. Examples include synchronisation, time warp and local perception filtering;
2. **Information Management Techniques** – These all reduce the amount of data that has to be managed by the network. Examples include Relevance Filtering and Compression;
3. **System Architecture Techniques** – These all aim to improve the efficiency of processing and disseminating data. Examples include protocols and network architectures.

This rest of this paper will focus on the first of these three categories and the other two categories will be discussed in part II of the survey. It should be noted that the techniques are not mutually exclusive and the DIA designer can mix techniques across categories to suit each particular application.

Time Management Techniques

Time is a fundamental attribute of all Distributed Interactive Applications, although the exact meaning of time may vary from one application to another. There are two dominant concepts of time for distributed applications:

1. **Absolute time** or **Wall Clock Time**: the time is based on the concept of a periodic clock, which is synchronized to Coordinated Universal Time (UTC) across the DIA. For example in the Internet this time is maintained by a network of servers that can communicate precise UTC time using the Network Time Protocol (NTP) (Mills 1991; Cox, Luijff, van Kampen and Ripley 1996);
2. **Virtual time** or **Causal Time**: the time in this case is based on a logical, loosely synchronized clock. Time is seen as a sequence of ordered events and stands still if no new events occur (Lamport 1978).

The link between time and consistency in DIAs is important. In a perfectly consistent DIA all participants would be rendered the same global state at the same absolute time. Network latency ensures that this cannot be achieved and inconsistencies therefore occur.

The concept of time within a DIA depends on whether the DIA is *continuous* or *discrete* (Roberts et al. 1997; Mauve et al. 2001; Zhou et al. 2001) and the events that are driving the state changes of the DIA. Events may be *deterministic* (e.g. a bouncing ball) or *non-deterministic* (e.g. an action by a user or network jitter). This categorisation of events allows more efficient use of bandwidth by transmitting non-deterministic events to other users and allowing deterministic events to be computed locally (Roberts et al. 1995). In a discrete DIA the state changes in response to non-deterministic user events. The correct order of events is therefore important but the exact time of each event is not. In this case an asynchronous model of time such as the

logical clock is sufficient. In a continuous DIA the state changes as a result of non-deterministic events caused by a user or by the passing of time, and so wall clock synchronization is important. In this case a correct order of all events must be maintained and the resulting state of the system must be the same as if all the events had been executed in the correct order at the time identified by their timestamps.

The foundations of time within DIAs were laid by Lamport (1978) who examined global time and the ordering of events in distributed systems. The underlying assumption is that events are related, that local events can be ordered sequentially and that the future cannot influence the past. This is referred to as the *causality relation (before-after relation)* and formed the basis for the concept of a logical clock. In such a clock the notion of time is based on the order in which local events occur. Time is therefore asynchronous among all participating nodes in the DIA. Virtual time differs from real time in that it doesn't flow and it stands still if no events occur.

Lamport's concept of scalar time as a totally ordered sequence of causal events is not always correct. With scalar time, all processes share the same global logical clock. This means that two or more events generated by different processes can have an identical timestamp, with the possible loss of causal dependency information (Raynal and Singhal 1996). For example, if a process generates an event based on a set of events it has received, but that new event is received at another node that hasn't received the same set of prior events, then a causality violation occurs. To solve this problem, a concept of time based on vector timestamps was developed independently by Mattern (1988) and Fidge (1988). With vector clocks, the time domain is represented by a set of vectors, one for the causal time of each process. This set of vectors represents the global time and each node maintains a vector clock for all processes. In a DIA the problem is then to disseminate the local vector clock to all other nodes. Various algorithms exist to do this. For example, the vector can get piggybacked on data packets.

For continuous DIAs, there is an additional constraint – it must appear as if the events have been executed at the correct point in time. In this case absolute time is needed and events are defined as a function of a wall clock, independently of the rendering frame rate. Various mechanisms exist for managing the clocks in the participating nodes. All clocks may be completely synchronized by using the Network Time Protocol (Mills 1991) or GPS (Cox et al. 1996). Lui (2001) derived an optimal synchronization interval based on subgraph communication algorithms so that every participating node has a consistent view of the virtual world.

Roberts and Sharkey (1997) combined the concepts of absolute and virtual time into a *sufficiently causal time stamp*. In doing so they modified the causality (before-after) relation, allowing the application to decide when to apply the causality relation to events. Some events may thus be ignored if a later event is absolute and the transitory state of no consequence. The ordering of events is termed *sufficient causal ordering* or *partial causal ordering* (Roberts et al. 1997).

The following sections describe the approaches that have been taken to maintain consistency in DIAs using both logical and wall clocks.

Lockstep Synchronization

This is the most basic way to ensure consistency and, being a conservative algorithm, avoids roll-back at the expense of system response (Funkhouser 1995). It operates by preventing the generation of out-of-order events and achieves this by preventing any participating node from advancing its simulation clock until all other nodes have acknowledged that they have completed their computation. This scheme guarantees consistency but does not guarantee real-time

consistency, as network latency results in acknowledgement delays and the DIA clock interval is not constant.

Imposed Global Consistency

Techniques that impose delays rely on delaying the execution or rendering of both local and remote events until an upper latency bound is reached. In this way the state information generated at remote nodes has sufficient time to be incorporated into the local computation of the global DIA state. The objective is to maintain synchronized global consistency at the expense of reduced responsiveness. This technique is a form of buffering and includes bucket synchronisation and local lag.

In bucket synchronisation, time is assumed to be synchronized and is divided into intervals or buckets of length T . The bucket frequency is thus $1/T$. All events are assigned to a bucket. The local view of the global state is then calculated using all the events generated locally and by remote nodes during the time interval $[i-D, i-D+T]$, where i is the current time interval and D is an added delay to compensate for network latency (Gautier et al. 1999). Events are therefore delayed for a length of time that should prevent any misorderings and hence avoid any rollbacks. The bucket synchronization algorithm, as applied to the distributed multiplayer game MiMaze, has been shown to provide acceptable consistency even in high latency situations (Gautier and Diot. 1998; Diot et al. 1999).

The concept of local lag is related to the work of Cristian (1985). Events initiated by the local user are not executed immediately but are delayed for a length of time dependent on network latency, thus reducing the responsiveness of the system. Each event is time stamped at a time later than the time the event actually occurred. An upper bound on the network delay is assumed so that events can be guaranteed to occur after a certain time. An event prologue can be added

between the actual local event and its local occurrence, which appears natural to the local user. The trade-off between short-term global inconsistencies and responsiveness when using local lag in replicated continuous applications was examined in (Mauve 2000). A minimal value for local lag was based on the maximum of the average network delay between two users, whereas the highest acceptable response time was based on human perception and application type, ranging from 100ms for drag operations to 400ms for click operations. The determination of an optimum local lag value was left to future work.

Delayed Global Consistency

In contrast to imposed global consistency, here the objective is to maintain asynchronous global consistency. Users will perceive the same consistent world but at different times.

Work by Qin on distributed whiteboards (Qin 2002) is based on the fact that it is not necessary that users have the same state at exactly the same time. Different nodes can share the same view of the objects but at a time shift. Each event is therefore time stamped so that the global state can be reconstructed locally, albeit at a later time. This shift is specified at the design stage of the DIA. Qin refers to this as delayed consistency.

A similar idea was proposed by Sharkey et al. (1998). They employed the concepts of relativity to DIAs to warp time and limit the maximum speed of objects within the virtual world. Each user is at the origin of a temporal-spatial framework and all other users are at a relative temporal-spatial position. This allows causal surfaces and volumes to be defined and associated critical velocities for virtual objects to be determined. The causal volumes define the regions in which past events can affect present events and present events can affect future events. The critical velocity determines the maximum speed of information transmission between two users of the DIA, hence delimiting the causal boundaries to produce causal surfaces. Based on these

causal surfaces, a $3^{1/2}D$ perception filter is proposed to introduce a continuously differentiable delay contour over the virtual environment. Each local user perceives other (remote) users interacting smoothly and dynamically with their local environment in real-time, but delayed by the particular latency between them and the local user. As users get closer and begin to interact, the delays between them are interpolated to zero. The causal volumes and the perception filter may potentially cause distortions in virtual time-space, leading to objects having apparent stiffness. Sharkey refers to this as *distributed stiffness*. In effect, a new set of consistent physics is established where users can understand delays and inconsistencies within the context of a relativistic physics paradigm, leading to constraints on time and space within the virtual world and as a result a new form of virtual physics fidelity.

Time Warp

The Time Warp mechanism, first proposed by Jefferson (1985; Jefferson 1989), is an optimistic synchronization mechanism that relies on general lookahead-rollback to maintain consistency in the DIA (Fujimoto 1990). It operates by executing each message as soon as it arrives. When a message arrives that has a time stamp earlier than a message already executed, it undoes all the events back to that message (rollback) and starts execution again from that message. It must also send messages to undo any incorrect output messages that were communicated to other nodes while it was in an incorrect state, a process known as rollback propagation (Lin and Lazowska 1991). To assist in rollback a snapshot of each state is taken after execution of each message. Jefferson extended the Time Warp mechanism to include optimal memory management and introduced the cancelback and the fossil collection protocols to facilitate this (Jefferson 1990). Jefferson's Time Warp mechanism was described using a formal model by Leivent and Watro (1993).

The Breathing Time Warp algorithm is a variation of the Time Warp algorithm. This limits the amount of execution that can be done optimistically by restricting the rollback time to events that occur within an event horizon (Steinman 1993).

Another optimistic variation on the Time Warp algorithm called Trailing State Synchronization was proposed in (Cronin et al. 2002) for distributed first-person shooter games. This requires less memory and processor time compared to time warp. A complete time-delayed copy of the game state is maintained in parallel with the main game. This copy has more time to reorder events and is a more accurate version of the game state, although delayed in time. This trailing state is used to detect inconsistencies. When an inconsistency is detected in the actual game, the correct state can then be directly copied from the trailing state.

Predictive Time Management

In predictive time management future events are pre-empted and both the event and its expected time of occurrence are transmitted before they actually occur. This is an optimistic technique for maintaining DIA consistency. The concept of predictive time management in DIAs was first proposed by Roberts et al. (1995) as part of the PaRADE system. A key motivation of their work was the issue of maintaining consistency among collaborating participants in the face of inherent network latency. Their prediction depends on knowledge of delays. They describe a study of Round Trip Time (RTT) values to estimate delays (Roberts et al. 1995). In PaRADE they use a sample set of message RTTs together with a robust line fitting algorithm to compute network latency between nodes. This technique relies on messages being time stamped, and thus a global wall clock is required. The prediction of events is limited to collision prediction (Sandoz, Sharkey and Roberts 1996) or predictions based on heuristics such as knowledge of intent, interest group or spatial proximity (Roberts et al. 1997). This is because sudden non-

deterministic events cannot be predicted and deterministic events are computed locally as the effects of the event can be described mathematically. Although the occurrence of non-deterministic events, initiated by humans-in-the-loop, are unpredictable certain long-term behaviours and strategies can be anticipated, and the user can be locked into a course of action in these cases (Delaney, Ward and Mc Loone 2003).

Events that are predicted locally are time-stamped and may be sent to other nodes where they are executed at the appropriate time. When using prediction it must be possible to predict further ahead than the network delay, so that messages sent to other participants arrive before the event occurs locally. The execution of local events can also be delayed so that the same event will be executed at the same time across the entire DIA. In this case additional filler events, known as *event prologues*, can be added between the initiation and the occurrence of local events to compensate for the local reduced responsiveness (Worthington et al. 2000). If the prediction is incorrect, roll back strategies are needed to return the DIA to a consistent state. To minimise the use of state roll back in PaRADE events are sent just-in-time by using knowledge of network delays.

Concurrency

The existence of latency means that if high performance interaction is required within the DIA, then the application state will probably be replicated at each participating node (Greenberg and Marwood 1994; Snowdon et al. 1996; Sun and Chen 2002). However, the replication of states increases the possibility of inconsistencies and this is further exacerbated by network latency due to geographical separation. These inconsistencies can generally be corrected by using an appropriate synchronisation algorithm such as time warp. However, when inconsistencies result from conflicts over ownership of objects, specific concurrency control

mechanisms are required. Such mechanisms may be optimistic (continue until a conflict is found then roll back) or pessimistic (check for conflict before allowing event). Concurrency is an important aspect of consistency and it is highly susceptible to network latency, especially when closely-coupled interaction is involved. Margery et al. (1999) define interaction as truly concurrent object manipulation, such as two or more users manipulating a shared object. Interaction may also be sequential, with different users interacting with the same, shared object but at different times. In DIAs, the issue of two or more users trying to access the same object in the DIA at the same time, either to perform some joint task or to obtain sole possession, must therefore be addressed.

Traditional concurrency control mechanisms used in databases such as locking and serialization are not suited to applications that require a fast response (Barghouti and Kaiser 1991). In groupware systems, a concurrency control mechanism known as Operational Transform is used (Ellis et al. 1989; Sun and Ellis 1998). This is an optimistic mechanism that replicates documents at each node and allows local operations to be performed immediately. Concurrent remote operations are transformed before execution so that both local and remote editing results are preserved. To resolve conflicting events in a replicated DIA while minimizing latency and maximising response time, Roberts et al. (1995) proposed a prediction-based concurrency control scheme. For one object to interact with another object in a DIA, the object initialising the interaction must possess a key (or token) to obtain ownership of the other object. In a DIA there is therefore a delay in requesting and receiving this key as it must be sent across the network, which makes it difficult to interact with the object on the first attempt. To combat this delay, Roberts et al. used predictive time management to request the key before the interaction actually takes place (Roberts et al. 1997). In their scheme the key unlocks a hierarchy

of related objects rather than just a single object. They also propose a number of techniques to avoid unnecessary passing of keys due to poor predictions.

In the prediction-based concurrency control scheme, ownership requests for the key are multicast to all potential owners of the object (Roberts et al. 1997). The process associated with the entity gathers all requests and predicts who the next owner will be. Lee et al. (2000) extended this concurrency control scheme by employing entity-based multicast and AOI. The AOI defines a spatial distance around the entity that currently possesses the key. All other entities in the AOI are potential future owners of the key. As entities enter/leave this AOI they join/leave the multicast group associated with the entity. The next owner of the key is then predicted from among all entities in the multicast group based on direction and predicted collision time.

VII. CONCLUDING REMARKS

This paper is part I of a two-part paper that examines the issues of latency and consistency in the context of Distributed Interactive Applications. Furthermore it describes various techniques and mechanisms that researchers have used to improve consistency and reduce latency effects. A classification for these mechanisms was proposed, based on three distinct categories – time management, information management and system architecture. Part I of this paper explored the category of time management and the various forms that this has taken in much of the research over recent decades. The choice of technique depends on the application area and, in particular, on the type of interaction required by participants. Each technique has its advantages and disadvantages. Delaying the execution of events, rolling back events or warping the physical environment all impact the experience of users participating in the DIA. It is therefore in the interest of the designers of such applications to balance consistency, responsiveness and fidelity

in the face of network latency. Part II of this paper continues with an examination of the two remaining categories of consistency maintenance mechanisms.

ACKNOWLEDGEMENT

This material is based upon works supported by Enterprise Ireland under grants no. SC/2002/129/ and Ircset/SC/04/CS0289.

REFERENCES

Barghouti, N. and G. Kaiser (1991). "Concurrency control in advanced database applications." *ACM Computing Surveys* **23**(3): pp. 269-317.

Barrus, J. W., R. C. Waters and D. B. Anderson (1996). "Locales: Supporting Large Multiuser Virtual Environments." *IEEE Computer Graphics and Applications* **16**(6): pp. 50-57.

Bhola, S., G. Banavar and M. Ahamad (1998). Responsiveness and Consistency Tradeoffs in Interactive Groupware. *Computer Supported Cooperative Work (CSCW '98)*, ACM, pp. 79-88.

Blakowski, G. and R. Steinmetz (1996). "A Media Synchronization Survey: Reference Model, Specification, and Case Studies." *IEEE Journal on selected areas in Communications* **14**(1): pp. 5-35.

Blanchard, T. D. and T. Lake (1997). *A Lightweight RTI Prototype with Optimistic Publication*. Spring Simulation Interoperability Workshop, Orlando, Florida, pp. 551-560.

Blau, B., C. E. Hughes, J. M. Moshell and C. Lisle (1992). *Networked Virtual Environments*. Symposium on Interactive 3D Graphics, Cambridge, Massachusetts, United States, ACM, pp. 157-160.

Blow, J. (1998). "A look at Latency in Networked Games." *Game Developer* **5**(7): pp. 28-40.

Bouras, C., G. Hornig, V. Triantafillou and T. Tsiatsos (2001). Architectures supporting e-learning through collaborative Virtual Environments: the case of INVITE. International Conference on Advanced Learning Technologies, Madison, Wisconsin, USA, pp. 13-16.

Bowman, D. A., D. B. Johnson and L. F. Hodges (2001). "Testbed Evaluation of Virtual Environment Interaction Techniques." Presence: Teleoperators and Virtual Environments **10**(1): pp. 75-95.

Broll, W. (1997). Distributed Virtual Reality for Everyone - a framework for networked VR on the internet. Virtual Reality Annual International Symposium (VRAIS '97), Albuquerque, New Mexico, IEEE, pp. 121-128.

Calvin, J., A. Dickens, B. Gaines, P. Metzger, D. Miller and D. Owen (1993). The Simnet Virtual World Architecture. Virtual Reality Annual International Symposium (VRAIS '93), Seattle, WA, USA, IEEE, pp. 450-455.

Capps, M., P. McDowell and M. Zyda (2001). "A Future for Entertainment-Defense Research Collaboration." IEEE Computer Graphics and Applications **21**(1): pp. 37-43.

Capps, M., D. McGregor, D. P. Brutzman and M. Zyda (2000). "NPSNET-V. A new beginning for dynamically extensible virtual environments." IEEE Computer Graphics and Applications **20**(5): pp. 12-15.

Capps, M. V. (2000). The QUICK Framework for Task-specific Asset Prioritization in Distributed Virtual Environments. Virtual Reality (VR '00), New Brunswick, NJ, USA, IEEE Computer Society, pp. 140-153.

Cheshire, S. (1996). Latency and the quest for interactivity. <http://www.stuartcheshire.org/>: A White paper commissioned by Volpe Welty Asset Management, L.L.C., for the synchronous person-to-person interactive computing environments meeting.

Churchill, E. F., D. N. Snowdon and A. J. Munro (2001). Collaborative Virtual Environments - Digital places and spaces for interaction. London, UK, Springer-Verlag.

Cox, A., E. Luijck, R. van Kampen and R. Ripley (1996). Time Synchronisation experiments. 14th Workshop on Standards for the Interoperability of Distributed Simulations, Florida, USA, pp. 599-613.

Cristian, F., H. Ahali, R. Stron and D. Dolev (1985). Atomic Broadcast: From Simple Message Diffusion to Byzantine Agreement. 15th International Symposium on Fault-tolerant Computing, Ann Arbor, MI, USA, IEEE Computer Society Press, pp. 200-206.

Cronin, E., B. Filstrup, A. R. Kurc and S. Jamin (2002). An Efficient Synchronization mechanism for Mirrored Game Architectures. First Workshop on Network and System Support for Games, Braunschweig, Germany, pp. 66-73.

de Oliveira, J. C. and N. D. Georganas (2002). VELVET: An Adaptive Hybrid Architecture for Very Large Virtual Environments. IEEE International conference on Communications, pp. 2491-2495.

de Oliveira, J. C., S. Shirmohammadi and N. D. Georganas (1999). Collaborative Virtual Environment Standards: A performance Evaluation. 3rd International Workshop on Distributed Simulation and Real Time Applications (DS-RT '99), College Park, Maryland, IEEE, pp. 14-21.

Delaney, D., T. Ward and S. Mc Loone (2003). Reducing Update Packets in Distributed Interactive Applications using a Hybrid Model. 16th International Conference on Parallel and Distributed Computing Systems, Reno, USA, pp. 417-422.

Diot, C. and L. Gautier (1999). "A Distributed Architecture for Multiplayer Interactive Applications on the internet." IEEE Network **13**(4): pp. 6-15.

Dourish, P. (1995). The Parting of Ways: Divergence, Data Management and Collaborative Work. Fourth European Conference on Computer Supported Cooperative Work (ECSCW '95), Stockholm, Kluwer Academic Publishers, Dordrecht, pp. 215-230.

Dutta-Roy, A. (2000). "The cost of quality in internet-style networks." IEEE Spectrum **37**(9): pp. 57-62.

Ellis, C. A. and S. J. Gibbs (1989). Concurrency control in groupware systems. Conference on Management of Data, Portland, OR, USA, ACM SIGMOD, pp. 399-407.

Ellis, S. R. (1994). "What are Virtual Environments?" IEEE Computer Graphics and Applications **14**(1): pp. 17-22.

Fidge, C. (1988). Timestamp in Message Passing Systems that Preserves Partial Ordering. 11th Australian Computing Conference, pp. 56-66.

Frecon, E. (2004). Dive on the Internet. Swedish Institute of Computer Science. PhD. Goteborg, IT University of Goteborg.

Frécon, E. (2003). DIVE: A generic tool for the development of virtual environments. 7th International Conference on Telecommunications (ConTEL 2003), Zagreb, Croatia, pp. 345-352.

Frécon, E. (2004). "DIVE: Communication architecture and programming model." IEEE Communications Magazine **42**(4): pp. 34-40.

Frécon, E. and M. Stenius (1998). "DIVE: A Scalable Network Architecture for Distributed Virtual Environments." Distributed systems Engineering Journal **5**: pp. 91-100.

Fujimoto, R. M. (1990). "Parallel Discrete Event Simulation." Communications of the ACM **33**(10): pp. 30-53.

Fujimoto, R. M. (2001). Parallel and Distributed Systems. Winter Simulation Conference, pp. 147-157.

Funkhouser, T. A. (1995). RING: A client-server system for multiuser virtual environments. Symposium on Interactive 3D Graphics, Monterey, California, United States, ACM SIGGRAPH, pp. 85-92.

Garcia, P., O. Montala, C. Pailot, R. Rallo and A. G. Skarmeta (2002). MOVE: Component Groupware Foundations for Collaborative Virtual Environments. 4th International Conference on Collaborative Virtual Environments (CVE02), Bonn, Germany, pp. 55-62.

Gautier, L., C. Diot and J. Kurose (1999). End-to-End Transmission Control Mechanisms for Multiparty Interactive Applications on the Internet. 18th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '99), New York, NY, USA, IEEE, pp. 1470-1479.

Gautier, L. and C. Diot. (1998). Design and evaluation of MiMaze, a multi-player game on the Internet. International Conference on Multimedia Computing and Systems (ICMCS '98), Austin, TX, IEEE, pp. 233-236.

Greenberg, S. and D. Marwood (1994). Real Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface. Computer Supported Cooperative Work (CSCW '94), North Carolina, ACM Press, pp. 207-218.

Greenhalgh, C. and S. Benford (1995). "MASSIVE: a collaborative virtual environment for teleconferencing." ACM Transactions on Computer-Human Interaction (TOCHI) 2(3): pp. 239-261.

Greenhalgh, C., J. Purbick and D. Snowdon (2000). Inside MASSIVE-3: Flexible Support for Data Consistency and World Structuring. 3rd International Conference on Collaborative Virtual Environments, San Francisco, California, ACM Press, pp. 119-127.

Greif, I. (1988). Computer-Supported Cooperative Work : A Book of Readings. San Mateo, California, Morgan Kaufmann.

Gutwin, C. (2001). The Effects of Network Delays on Group Work in Real-Time Groupware. Seventh European Conference on Computer-supported Cooperative Work (ECSCW '01), Boon, Germany, Kluwer Academic Publishers, Dordrecht, pp. 299-318.

Hecker, C. and Z. B. Simpson (2001). "Dead Reckoning a.k.a. Motion Prediction." *Game Developer* **8**(2): pp. 10.

Henderson, T. N. H. (2003). The effects of relative delay in networked games. Computer Science. PhD. London, University College London: 150.

Holbrook, H. W., S. K. Singhal and D. R. Cheriton (1995). "Log-based Receiver-reliable Multicast for distributed Interactive Simulation." *ACM SIGCOMM Computer Communications Review* **25**(4): pp. 328-334.

IEEE (1993). IEEE Standard for Distributed Interactive Simulation - Application Protocols. IEEE Std 1278-1993. C. Society. New York, IEEE.

IEEE (1995). IEEE Standard for Distributed Interactive Simulation - Application Protocols. IEEE Std 1278.1-1995 (Revision of IEEE Std 1278-1993), IEEE.

Jefferson, D. R. (1985). "Virtual time." *ACM Transactions on Programming Languages and Systems* **7**(3): pp. 404-425.

Jefferson, D. R. (1989). "Virtual Time." *ACM Transactions on Programming Languages and Systems* **7**(1): pp. 404-425.

Jefferson, D. R. (1990). Virtual Time II: Storage Management in Distributed Simulation. 9th Annual Symposium on Principles of Distributed Computing, ACM, pp. 75-89.

Jehaes, T., D. De Vleeschauwer, T. Coppens, B. Van Doorselaer, E. Deckers, W. Naudts, K. Spruyt and R. Smets (2003). Access Network Delay in Networked Games. 2nd Workshop on Network and System Support for Networked Games (NetGames '03), Redwood City, California, ACM, pp. 63-71.

Joslin, C., T. Di Giacomo and N. magnenat-Thalmann (2004). Collaborative Virtual Environments: From birth to Standardization. IEEE Communications Magazine. **44**: 28-33.

Joslin, C., T. Molet, N. Thalmann, J. Esmerado, D. Thalmann, I. Palmer, N. Chilton and R. Earnshaw (2001). "Sharing Attractions on the Net with VPark." IEEE Computer Graphics and Applications **21**(1): pp. 61-71.

Kelly, T. (1997). Dr. T: A Communications Infrastructure for Distributed Interactive Applications. Paper available on the web.

Krumm-Heller, A. and S. Taylor (2000). Using determinism to improve the accuracy of dead reckoning algorithms. SimTecT2000 - The Simulation Technology and Training Conference, Sydney, Australia, pp.

Lamport, L. (1978). "Time, clocks and the ordering of events in a distributed system." Communications of the ACM **21**(7): pp. 558-565.

Lee, D., M. Lim and S. Han (2002). ATLAS - A Scalable Network Framework for Distributed Virtual Environments. 4th International Conference on Collaborative Virtual Environments (CVE02), Bonn, Germany, pp. 47-54.

Lee, D., J. Yang, H. Y. Yoon, C. Yu and S. J. Hyun (2000). Entity-Centric Scalable Concurrency Control for Distributed Interactive Applications. International Performance, Computing, and Communications Conference (IPCCC '00), Phoenix, Arizona, USA, IEEE, pp. 544-550.

Leigh, J., O. Yu, D. Schonfeld, R. Ansari, E. He, A. Nayak, J. Ge, N. Krishnaprasad, K. Park, Y.-j. Cho, L. Hu, R. Fang, A. Verlo, L. Winkler and T. A. DeFanti (2001). Adaptive Networking for Tele-Immersion. 5th Immersive Projection Technology/ 7th Eurographics Virtual Environments Workshop (IPT/EGVE), Stuttgart, Germany, pp. May 16-18.

- Levient, J. I. and R. J. Watro (1993). "Mathematical Foundations for Time Warp Systems." ACM Transactions on Programming Languages and Systems **15**(5): pp. 771-794.
- Lin, Y.-B. and E. D. Lazowska (1991). "A Study of Time Warp Rollback Mechanisms." ACM Transactions on Modelling and Computer Simulation **1**(1): pp. 51-72.
- Lui, J. C. S. (2001). "Constructing Communication Subgraphs and Deriving an Optimal Synchronization Interval for Distributed Virtual Environment Systems." IEEE Transactions on Knowledge and Data Engineering **13**(5): pp. 778-792.
- MacColl, I., D. Millard, R. C. and A. Steed (2002). Shared Visiting in EQUATOR City. 4th International Conference on Collaborative Virtual Environments (CVE02), Bonn, Germany, pp. 88-94.
- Macedonia, M. and M. Zyda (1997). "A Taxonomy for Networked Virtual Environments." IEEE Multimedia **4**(1): pp. 48-56.
- Macedonia, M. R. (1995). A Network Software Architecture for Large Scale Virtual Environments. Department of Computer Science. PHD. Monterey, CA 93943-5000, Naval Postgraduate School: 232.
- Macedonia, M. R., M. J. Zyda, D. R. Pratt, P. T. Barham and S. Zeswitz (1994). "NPSNET: A Network Software Architecture for Large Scale Virtual Environments." Presence **3**(4): pp. 265-287.
- Manninen, T. (2000). Interaction in networked Virtual Environments as Communicative Action: Social theory and Multi-player games. 6th International Workshop on Groupware (CRIWG2000), Madeira, Portugal, IEEE Computer Society Press, pp. 154-157.
- Margery, D., B. Arnaldi and N. Plouzeau (1999). A General Framework for Cooperative Manipulation in Virtual Environments. Eurographics Workshop, Vienna, Austria, Springer, pp. 169-178.

Martin, H. S., A. J. McGregor and J. G. Cleary (2000). Analysis of Internet Delay Times. 1st Passive and Active Measurements Workshop (PAM '00), Hamilton, New Zealand, pp. 141-149.

Mattern, F. (1988). Virtual Time and Global States of Distributed Systems. International Workshop on Parallel and Distributed Algorithms, France, Elsevier Science Publishers B.V. (North-Holland), pp. 215-226.

Mauve, M. (2000). Consistency in Replicated Continuous Interactive Media. Computer Supported Cooperative Work (CSCW '00), Philadelphia, PA, ACM, pp. 181-190.

Mauve, M. (2000). Distributed Interactive Media. Berlin.

Mauve, M., V. Hilt, C. Kuhmunch and W. Effelsberg (2001). "RTP/I - Towards a Common Application Level Protocol for Distributed Interactive Media." IEEE Transactions on Multimedia **3**(1): pp. 152-161.

Mauve, M., J. Vogel, V. Hilt and W. Effelsberg (2002). "Local-lag and Timewarp: Providing Consistency for Replicated Continuous Applications." To appear in: IEEE Transactions on Multimedia, 2002 (not there 21.02.2003): pp.

Mauve, M., J. Vogel, V. Hilt and W. Effelsberg (2004). "Local-lag and Timewarp: Providing Consistency for Replicated Continuous Applications." IEEE Transactions on Multimedia **6**(1): pp. 47-57.

McCoy, A., D. Delaney and T. Ward (2003). "Game-state fidelity across distributed interactive games." ACM Crossroads, Networking Issue 9.4, Summer: pp. 4-9.

Meehan, M., S. Razzaque, M. C. Whitton and F. P. j. Brooks (2003). Effect of Latency on Presence in Stressful Virtual Environments. Virtual Reality (VR '03), IEEE Computer Society, pp. 141-148.

Mills, D. L. (1991). "Internet Time Synchronization: The Network Time Protocol." IEEE Transactions on Communications **39**(10): pp. 1482-1493.

Natrajan, A. and P. F. Reynolds Jr (1999). Resolving Concurrent Interactions. 3rd International Workshop on Distributed Simulation and Real Time Applications (DS-RT '99), College Park, Maryland, IEEE, pp. 85-92.

Park, K. S. and R. V. Kenyon (1999). Effects of Network Characteristics on Human Performance in a Collaborative Virtual Environment. Virtual Reality (VR '99), Houston, Texas, IEEE Computer Society, pp. 104-111.

Pullen, J. M. and D. C. Wood (1995). "Network Technology and DIS." Proceedings of the IEEE **83**(83): pp. 1156-1167.

Qin, X. (2002). "Delayed Consistency Model for Distributed Interactive Systems with Real-time Continuous Media." Journal of Software **13**(6): pp. 1029-1039.

Radhakrishnan, R. and P. A. Wilsey (1999). Ruminations on the implications of Multi-Resolution modelling on DIS/HLA. 3rd International Workshop on Distributed Simulation and Real Time Applications (DS-RT '99), College Park, Maryland, IEEE, pp. 101-108.

Raynal, M. and M. Singhal (1996). "Logical Time: Capturing Causality in Distributed Systems." Computer **29**(2): pp. 49-56.

Riva, G. and L. Gamberini (2001). "Virtual Reality in Telemedicine." Telemedicine Journal and e-Health **6**(3): pp. 327-340.

Roberts, D. J. (2004). Communication Infrastructures for Inhabited Information Spaces. Inhabited Information Spaces: Living With Your Data (Computer Supported Cooperative Work).

D. N. Snowdon, E. F. Churchill and E. Frecon. London, Springer-Verlag: 233-267.

Roberts, D. J. and P. M. Sharkey (1997). Maximising Concurrency and Scalability in a Consistent, Causal, Distributed Virtual Reality System, whilst minimising the effect of Network Delays. 6th Workshop on Enabling Technologies on Infrastructure for Collaborative Enterprises (WET-ICE '97), Cambridge, MA, IEEE, pp. 161-166.

Roberts, D. J. and P. M. Sharkey (1997). "Minimising the latency induced by consistency control within a large scale multi-user distributed virtual reality system." IEEE Conference on Systems, Man and Cybernetics **5**: pp. 4492-4497.

Roberts, D. J., P. M. Sharkey and P. D. Sandoz (1995). A Real-time Predictive Architecture for Distributed Virtual Reality. 1st Workshop on Simulation and Interaction in Virtual Environments, Iowa City, ACM SIGGRAPH, pp. 279-288.

Roza, M., J. Voogd, H. Jense and P. van Gool (1999). Fidelity Requirements Specification: A Process Oriented View. Fall Simulation Interoperability Workshop, Orlando, FL, pp. Paper 99F-SIW-032.

Sandoz, P. D., P. M. Sharkey and D. J. Roberts (1996). Collision Prediction of a Moving Object Within a Virtual World. Virtual Reality World '96, Stuttgart, Germany, pp.

Schwarz, R. and F. Mattern (1994). "Detecting Causal Relationships in Distributed Computations: In Search of the Holy Grail." Distributed Computing **7**(3): pp. 149-174.

Sharkey, P. M., M. D. Ryan and D. J. Roberts (1998). A local perception filter for Distributed Virtual Environments. Virtual Reality Annual International Symposium (VRAIS '98), Atlanta, Georgia, IEEE, pp. 242-249.

Shneiderman, B. (1984). "Response Time and Display Rate in Human Performance with Computers." Computing Surveys **16**(3): pp. 625-685.

Singhal, S. K. (1996). Effective Remote Modelling in Large-scale Distributed Simulation and visualisation environments. PhD Thesis. Stanford University.

Singhal, S. K. and M. Zyda (1999). Networked Virtual Environments. New York, ACM Press.

Smed, J., T. Kaukoranta and H. Hakonen (2001). Aspects of Networking in Multiplayer Computer Games. International Conference on Application and Development of Computer Games in the 21st Century, China, pp. 74-81.

Snowdon, D., C. Greenhalgh, S. Benford, A. Bullock and C. Brown (1996). "A Review of Distributed Architectures for Networked Virtual Reality." *Virtual Reality: Research, Development and Applications* **2**(1): pp. 155-175.

Steinman, J. S. (1993). Breathing Time Warp. 7th International Workshop on Parallel and Distributed Simulation (PADS93), pp. 109-118.

Stytz, M. R. (1996). "Distributed Virtual Environments." *IEEE Computer Graphics and Applications* **16**(3): pp. 19-31.

Sun, C. and D. Chen (2002). "Consistency Maintenance in real-time collaborative graphics editing systems." *ACM Transactions on Computer-Human Interaction (TOCHI)* **9**(1): pp. 1-41.

Sun, C. and C. A. Ellis (1998). Operational Transformation in real-time group editors: issues, algorithms and achievements. *Computer Supported Cooperative Work (CSCW '98)*, ACM Press, pp. 59-68.

Sun, C., X. Jia, Y. Zhang, Y. Yang and D. Chen (1998). "Achieving Convergence, Causality preservation and intention preservation in real-time cooperative editing systems." *ACM Transactions on Computer-Human Interaction (TOCHI)* **5**(1): pp. 63-108.

Tawfik, H. and T. Fernando (2001). A Simulation Environment for Construction Site Planning. 5th International Information Visualisation Conference, London, England, IEEE, pp. 199-204.

Tse-Au, E. S. H. and P. A. Morreale (2000). End-to-End QoS measurement: Analytic methodology of application response time vs. tunable latency in IP networks. *IEEE/IFIP Network operations and management symposium (NOMS '00)*, Honolulu, USA, pp. 129-142.

Vaghi, I., C. Greenhalgh and S. Benford (1999). Coping with inconsistency due to network delays in collaborative virtual environments. *Symposium on Virtual Reality Software and Technology (VSRT '99)*, London, England, ACM, pp. 42-49.

Vogel, J., M. Mauve, V. Hilt and W. Effelsberg. (2003). "Late Join Algorithms for Distributed Interactive Applications." *Multimedia Systems* **9**(4): pp. 327-336.

Waters, R. C. and J. W. Barrus (1997). "The rise of shared virtual environments." *IEEE Spectrum* **34**(3): pp. 20-25.

Wloka, M. (1995). "Lag in Multiprocessor VR." *Presence: Teleoperators and Virtual Environments* **4**(1): pp. 50-63.

Worthington, B. G. and D. J. Roberts (2000). Encapsulating Network Latency Compensators in VRML. *Virtual Worlds and Simulation Conference (VWSIM '00)*, San Diego, CA, USA, pp. Int. Conf. SCS Western Multi-conference.

Zeigler, B. P., H. Praehofer and T. G. Kim (2000). Theory of Modeling and Simulation. San Diego, CA, USA, Academic Press.

Zhou, S., W. Cai, B.-s. Lee and S. J. Turner (2003). "Time-space Consistency in Large Scale Distributed Virtual Environment." *ACM Transactions on Modeling and Computer Simulation* **14**(1): pp. 31-47.

Zhou, S. P., W. T. Cai, F. B. S. Lee and S. J. Turner (2001). Consistency in Distributed Interactive Applications. *European Simulation Interoperability Workshop 2001 (Euro-SIW 2001)*, London, England, pp. 01E-SIW-003.

Figure 1: Illustration of network nodes. B refers to bandwidth, τ refers to processing time and $\Delta\tau$ refers to propagation delay.

Size of Figure 1: Height: 6.00cm Width: 12.84cm