

Examining the Effects of Time-Space Measures on the Hybrid Strategy Model in Networked Virtual Environments

Damien Marshall, Séamus McLoone, Declan Delaney, and Tomás Ward

*Department of Electronic Engineering,
National University of Ireland, Maynooth,
Maynooth, County Kildare, IRELAND*

E-mail: damien.marshall@eeng.nuim.ie, seamus.mcloone@eeng.nuim.ie

Scalability is an important issue in the design of Networked Virtual Environments (NVEs). In order to achieve scalability, it is essential to minimise the network traffic required to maintain overall consistency in the NVE. A popular method of achieving this is via entity behaviour prediction mechanisms, such as dead reckoning and the Hybrid Strategy Model (HSM). Typically, the performance of such mechanisms is rated by the number of network packets they generate. However, it is also important that their impact on overall consistency is investigated. Absolute consistency is the degree to which different views of a NVE on remote hosts correspond. In previous work, it was shown that the use of a spatial threshold with dead reckoning can result in unbounded local absolute consistency. A solution that employed a time-space error threshold measure was shown to remedy this issue. In this paper, the scope of the time-space measure is extended to include the HSM. It is shown how the HSM can also result in unbounded local absolute inconsistency. A solution that once again incorporates the time-space threshold is examined. However, this approach results in a significant increase in network traffic. To resolve this, a novel extension to the HSM algorithm is presented, which is demonstrated to reduce network traffic, whilst still maintaining a low level of local absolute inconsistency.

Keywords – Networked Virtual Environments, Inconsistency, Time-Space Threshold, Dead Reckoning, Hybrid Strategy Model

I INTRODUCTION

Scalability is the ability of a Networked Virtual Environment (NVE) to deal with additional data [1,2]. This data may result from an increase in the number of participants, from an increase in the complexity of the tasks they are executing or from an increase in the level of interaction among participants in the virtual environment. One approach to increase the scalability of the NVE is to reduce the amount of network traffic transmitted during execution of the environment. A popular method for achieving this is entity prediction [3,4].

Entity prediction mechanisms operate by employing both a high and low fidelity model of entity behaviour at the home node of each entity. The home node continually compares the high fidelity model of the entity (its actual behaviour) with a low fidelity model of the same entity (which is also used at remote hosts as the predicted behaviour). The

home node transmits an update network packet when a significant level of inconsistency between the high and low fidelity model occurs. Examples of entity prediction mechanisms are dead reckoning and the Hybrid Strategy Model (HSM) approach [3,4].

Under the dead reckoning scheme, the low fidelity model employs entity dynamics to predict future behaviour in the short term. However, in addition to entity dynamics, the HSM also employs a priori information regarding entity behaviour in relation to an objective or goal in the environment in order to predict long-term user behaviour. Previous work demonstrated that the HSM approach can reduce network traffic in comparison to an approach based on dead reckoning alone [4].

In order to fully evaluate the impact of employing such mechanisms, it is necessary to examine both the inconsistency resulting from their use, as well as the reduction in network packets generated. Consistency is the degree to which

different views of a NVE on remote hosts correspond [5].

Normally, both dead reckoning and the HSM use spatial inconsistency measures to determine the generation and transmission of network packets. In previous work, it was demonstrated that using a purely spatial inconsistency threshold for dead reckoning can lead to unbounded absolute inconsistency [6]. A time-space threshold metric to replace the existing dead reckoning spatial error threshold was subsequently proposed. This threshold uses time-space inconsistency measures to determine when an update is required, and it was shown to result in bounded local absolute inconsistency but unbounded spatial inconsistency [7]. By combining both the spatial and the time-space metrics to form a hybrid threshold criteria, both spatial and absolute inconsistency were bounded.

In this paper, we investigate the use of the time-space threshold metric in the context of the Hybrid Strategy Model and perform a comparative analysis between the HSM and dead reckoning approaches. We identify a shortcoming in the HSM approach when the time-space metric is employed. It is experimentally demonstrated that a novel enhancement to the HSM involving “strategy shifting” results in the generation of fewer packets compared to the use of dead reckoning, whilst still placing a bound on absolute inconsistency.

The rest of the paper is organised as follows. Section 2 discusses the HSM and various measures of the different aspects of inconsistency. Section 3 explores the use of the time-space measure with the HSM, and exposes a weakness in the HSM approach when the two are combined. A novel solution to this problem is presented in Section 4, and subsequently verified in Section 5. The paper ends with some conclusions and suggestions for future work in Section 6.

II RELATED WORK

For the convenience of the reader, we will now discuss some related work, concerning the Hybrid Strategy Model, and different measures of inconsistency.

a) Hybrid Strategy Model

The Hybrid Strategy Model operates by combining a short-term model, such as dead reckoning, and a long-term predictive model, such as the strategy model of user behaviour [4]. The strategy model is a long-term approach that exploits the fact that, when given a set goal in a static environment, users tend to converge on a preferred trajectory or strategy to meet the goal. The high fidelity model is analysed, and the appropriate long-term strategy model is chosen based on this behaviour. The low fidelity model then follows the appropriate strategy model. As long as

the high fidelity model of user behaviour matches that of the low fidelity model to within a certain threshold, no network packets are transmitted.

Although this approach has been shown to perform better than dead reckoning in terms of network traffic, it is important that the performance of all such mechanisms be measured from the perspective of inconsistency.

b) Measures of Inconsistency

Several measures of the different factors of inconsistency have been proposed by various authors. Diot and Gautier use an inconsistency metric known as “drift distance” to measure the inconsistency of their MiMaze application [8]. Drift distance is the spatial difference between the high fidelity and low fidelity model position on the remote host at each time step. This measurement ignores the duration of the inconsistency. On the other hand, Lui proposes a measure which is based on time only [9]. Lui proposes that inconsistency is the difference between the start times of rendering of the same changes to the virtual environment at different nodes. This is known as “phase difference”. Another interesting factor of consistency is that of semantic inconsistency [10]. According to Correa et al., semantic consistency pertains to the meaning of the data present in the virtual environment. The authors measure semantic consistency in order to reduce network traffic that needs to be communicated amongst nodes with heterogeneous resources.

Zhou et al. provide a method of measuring absolute inconsistency, which combines temporal and spatial inconsistency in a single measure [11]. The measure of absolute inconsistency, Ω is given in Equation 1 below:

$$\Omega = \begin{cases} 0, & \text{if } |\Delta(t)| < \varepsilon \\ \int_{t_0}^{t_0+\tau} |\Delta(t)| dt, & \text{if } |\Delta(t)| \geq \varepsilon \end{cases} \quad (1)$$

where

- Δ is the difference between the position of a local object and its remote replication;
- $\Delta(t)$ is the above difference over a duration t ;
- t_0 is the start time of the inconsistency and τ is the duration;
- ε is the minimum perceivable distance.

In other words, inconsistency is measured as the area under the curve given by spatial inconsistency over time. Here, inconsistency has to be perceivable by the end user. Zhou et al. use this metric to evaluate the inconsistency arising from a Networked Virtual Environment simulation before its execution.

In previous work we demonstrated, using Zhou’s metric, that local absolute inconsistency can be unbounded when a spatial threshold is employed with dead reckoning [6]. This occurred because the high fidelity model could remain spatially inconsistent for extended periods of time. The

“inline” update, or an update within the bounds of a strategy model.

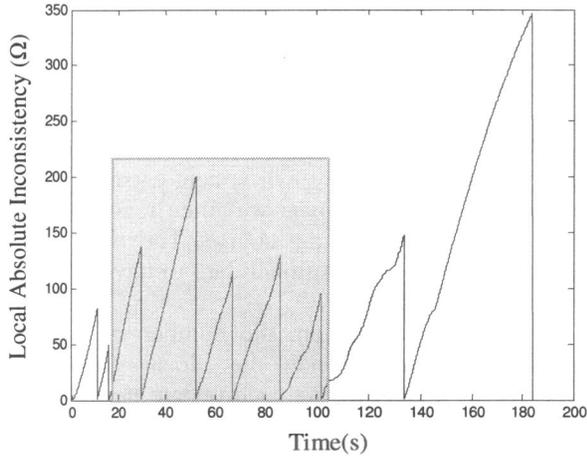


Figure 3. Local Absolute inconsistency for the user trial shown in Figure 2. The inconsistency relevant to Figure 2 is highlighted.

However, consider the local absolute inconsistency arising during the execution of the same simulation, shown in Figure 3. Here, the local absolute inconsistency arising during the whole simulation is plotted against time. The inconsistency for the section of the trial shown in Figure 2 is highlighted.

It can be seen that local absolute inconsistency is unstable, and potentially unbounded, in this example. Although a similar issue arises when dead reckoning is employed with a spatial threshold, the problem is exacerbated with the HSM, as the time between updates can be significantly longer. For example, in Figure 2, the time between the second and third update, and the third and fourth update is approximately 10 seconds and 20 seconds, respectively. The largest amount of local absolute inconsistency arises in the highlighted section of Figure 3 during these times.

a) Time-Space Threshold

As with dead reckoning, the problem of unbounded local absolute inconsistency can be remedied by incorporating the time-space measure as an error threshold, i.e. we employ a combination of the time-space and spatial thresholds.

To test this approach, the simulations were repeated in the test platform using the combined time-space/spatial threshold and the Hybrid Strategy Model. Again, the local absolute inconsistency and number of update packets were recorded for each simulation run.

As would be expected, local absolute inconsistency is then bounded to the error threshold value, as shown in Figure 4. Here, the absolute inconsistency arising from the simulation presented in Figure 2 is shown when a combined time-space/spatial threshold of 10 is employed.

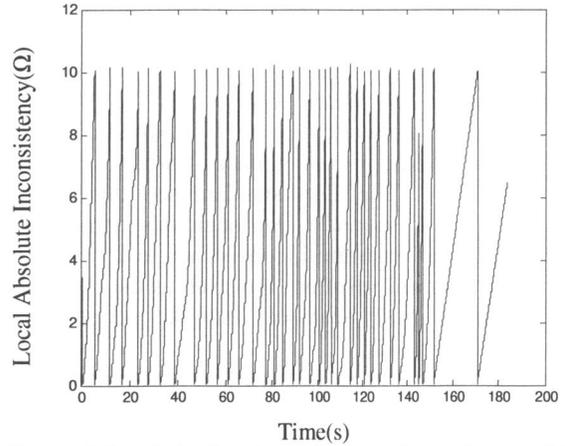


Figure 4. Local absolute inconsistency arising when combined time-space/spatial threshold of 10 is employed.

However, the performance of the HSM from the point of view of network packet update generation is adversely affected. Each peak in Figure 4 represents a point where an update is transmitted. It is evident that the number of peaks has increased dramatically in comparison to Figure 3. There is a clear trade off between improved local absolute consistency, and increasing the number of update packets. In order to analyse the extent of this issue, the “bot” simulations were repeated using dead reckoning and the combined time-space/spatial thresholds.

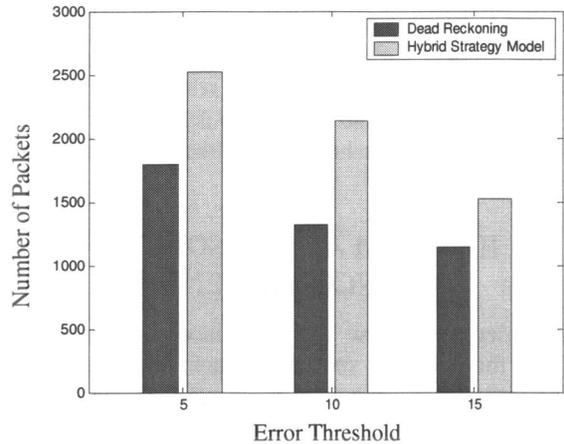


Figure 5. Comparison of Dead Reckoning and the Hybrid Strategy Model with combined time-space/spatial thresholds

Figure 5 shows the total number of simulated packets generated during both of the HSM and dead reckoning simulations, for combined time-space/spatial thresholds of 5, 10 and 15. In each case, the error threshold of both dead reckoning and the HSM is equivalent. The distance between parallel models was set to 1.6 times the error threshold value. A total of 3 parallel models were employed in each case.

These results demonstrate that dead reckoning performs significantly better than the HSM overall when the combined time-space/spatial threshold is employed. For an error threshold of 10, the HSM transmits approximately 61% more packets than dead reckoning. This contradicts previous work in

this area, which demonstrated that the HSM requires less update packets than dead reckoning.

Upon closer examination of the recorded data, it was found that the time-space threshold was being breached quickly while the high fidelity model remained within in the bounds of the strategy model. This caused a significant increase in the frequency of “in-line” updates, as demonstrated in Figure 6. This behaviour can be attributed to the manner in which the low fidelity model is represented using the long term strategy model.

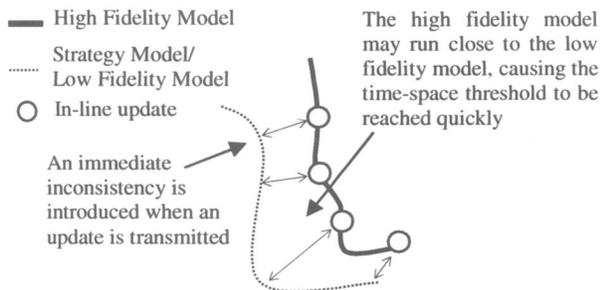


Figure 6. Using the time-space threshold with the Hybrid Strategy Model increases the generation of “in-line” updates

Firstly, when an update is generated, the low fidelity model is set to the closest position on the relevant long-term strategy model, introducing an immediate spatial inconsistency. Such a problem does not arise with dead reckoning, as the low fidelity model is positioned at the location of the high fidelity model where the update occurs.

Secondly, the high fidelity model and the long-term strategy model may run close to each other. Although the spatial error threshold is not exceeded, a spatial inconsistency exists over an extended period of time, resulting in the time-space threshold being reached quickly.

To alleviate this problem, we now propose a novel extension to the HSM mechanisms that allows for the use of the time-space threshold and its benefits, yet reduces the number of update packets in comparison to dead reckoning.

IV STRATEGY SHIFTING

In Section 3, it was highlighted how the increase in update generation caused by the use of the time-space threshold was related to the manner in which the low fidelity model was updated.

To remedy this issue, it is proposed here that when a network update is generated, the strategy model is shifted to the position where the update occurred – see Figure 7. The shifted model is parallel to the originating strategy model.

This approach has three main benefits. (1) The initial spatial inconsistency introduced when the low fidelity model position is set to a position on the strategy model is removed; (2) The shifted strategy model can accommodate those high fidelity models that run

parallel to the strategy model; (3) The time-space threshold can be employed, without the need for the constant generation of in-line network updates.

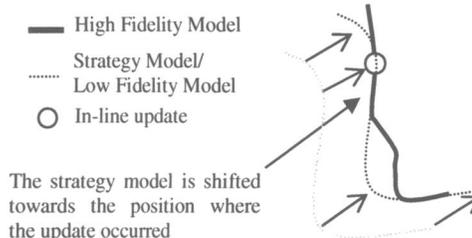


Figure 7. Shifting the strategy model to the position where the update occurs reduces the frequency of inline updates

In the next section, this new implementation of the HSM is verified via simulation and comparison with the Dead Reckoning approach. An analysis of the results is also presented.

V RESULTS AND ANALYSIS

The Torque experimental test platform was extended to include the strategy shifting mechanism with the Hybrid Strategy Model. The user data discussed in Section 2 was again employed to allow a “bot” entity to navigate the test environment, while employing the new HSM approach. The threshold value was varied as required. As only one model exists at any one time, there is no need to specify a distance between models, or the number of models required. If the entity moved to a distance of 30 Torque game units from the original strategy model, dead reckoning was employed. However, due to the nature of the course, all participants stayed within the bounds of the environment track, and dead reckoning was not utilised during the simulations.

Figure 8 shows the total number of simulated packets for all simulations. As before, combined time-space/spatial error thresholds of 5, 10 and 15 were employed.

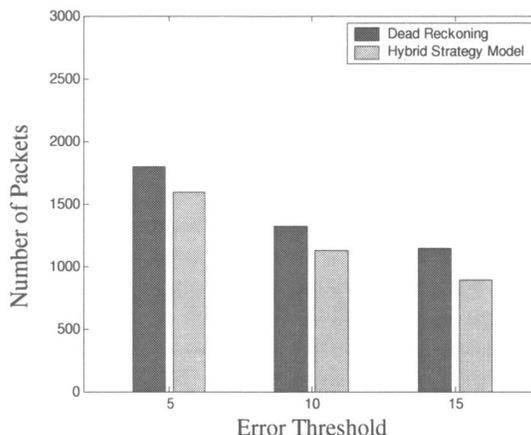


Figure 8. Comparison of Dead Reckoning and the Hybrid Strategy Model when shifted strategy models are employed

As is expected, the required number of network packets decreases as the threshold increases for both

the HSM and dead reckoning. The situation is reversed in comparison to the data presented in Figure 5. HSM now performs better than dead reckoning overall, while still incorporating the benefits of the time-space threshold.

These benefits are further demonstrated in Figure 9. Here, a magnified section of the low and high fidelity model from a single user trial is presented, along with an original strategy model, and the shifted strategy model. A combined spatial/time-space threshold of 10 was employed. When the first update occurs, the strategy model is shifted to a position on the high fidelity model. The low fidelity model follows this new path.

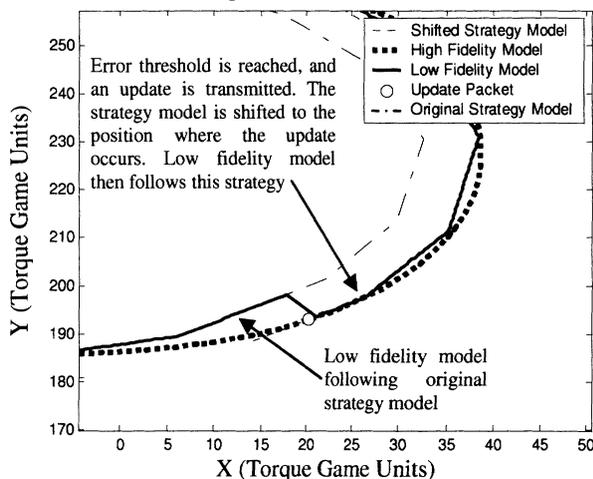


Figure 9. Magnified section of a simulation using the combined spatial/time-space threshold and strategy shifting.

Subsequently, the high fidelity model remains parallel to the low fidelity model, introducing a constant spatial inconsistency over time. As the combined is being employed, this situation is quickly detected. A further update is transmitted. The strategy model is again shifted to the high fidelity model. We can then see that the low fidelity model matches the high fidelity model closely, without any need for further updates in the short term.

VI CONCLUSIONS

In this paper, our work with the time-space threshold was extended to include the Hybrid Strategy Model (HSM). It was demonstrated that the use of a spatial threshold with the HSM results in poor local absolute consistency, similar to that of dead reckoning.

To remedy this, a combined spatial/time-space metric was incorporated as a threshold in HSM. Although this placed a bound on local absolute inconsistency, the number of update packets was found to increase significantly. The HSM then performed worse than dead reckoning. This increase is caused by the generation of extra "in-line" updates, or updates within the bounds of a strategy model itself.

A novel extension to the HSM was then proposed to combat this issue. By shifting the strategy model

to the position where the update occurred, the frequency of inline updates was reduced, but still allowed for the benefits provided from the use of the time-space threshold. This new approach performed better than dead reckoning.

Future work will examine methods of comparing spatial and time-space thresholds, by relating them from a psychoperceptual perspective.

VII ACKNOWLEDGEMENTS

This work is supported by the Enterprise Ireland under grant number IRCSET/SC/04/CS0289

VIII REFERENCES

- [1] Roberts, D. J. and P. M. Sharkey "Maximising Concurrency and Scalability in a Consistent, Causal, Distributed Virtual Reality System, whilst minimising the effect of Network Delays". *Proceedings of the 6th Workshop on Enabling Technologies Infrastructure for Collaborative Enterprises*, Cambridge, MA, IEEE, pp. 161-166, 1997
- [2] Bondi, A. B. "Characteristics of Scalability and Their Impact on Performance". *Second international workshop on Software and Performance*, Ottawa, Ontario, Canada, ACM Press, pp. 195 - 203, 2000
- [3] IEEE. *IEEE Standard for Distributed Interactive Simulation - Application Protocols IEEE Std 1278.1-1995* IEEE. (1995)
- [4] Delaney, D., T. Ward and S. McLoone "On Reducing Entity State Update Packets in Distributed Interactive Simulations using a Hybrid Model". *Proceeding of the 21st IASTED International Multi-conference on Applied Informatics, February 10-13, Innsbruck, Austria*, pp. 2003
- [5] Delaney, D., T. Ward and S. McLoone. "On consistency and network latency in distributed interactive applications: A survey - Part I." *Presence: Teleoperators and Virtual Environments* 15(2), April 2006
- [6] Roberts, D., D. Marshall, S. McLoone, D. Delaney, T. Ward and R. Aspin "Exploring the use of local inconsistency measures as thresholds for dead reckoning update packet generation". *Distributed Simulation and Real Time Applications*, Montreal, Canada, pp. 195 - 202, 2005
- [7] Marshall, D., D. Roberts, D. Delaney, S. McLoone and T. Ward "Dealing with the Effect of Path Curvature on Consistency of Dead Reckoned Paths in Networked Virtual Environments". *IEEE Virtual Reality Conference*, Virginia, USA, pp. 299-300, 2006
- [8] Diot, C. and L. Gautier. "A Distributed Architecture for Multiplayer Interactive Applications on the internet." *IEEE Network* 13(4):pp. 6-15, 1999
- [9] Lui, J. C. S. "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, 2001
- [10] Correa, C., I. Marsic and X. Sun. *Semantic Consistency Optimization in Heterogeneous Virtual Environments*, Technical Report CAIP-TR-267 Rutgers University, CAIP Center. (2002)
- [11] Zhou, S., W. Cai, F. B. S. Lee and S. J. Turner. "Time-space Consistency in Large Scale Distributed Virtual Environment." *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 14(1):pp. 31 - 47, 2004
- [12] Hook, D. V., J. Calvin and J. Smith "Data Consistency Mechanisms to Support Distributed Simulations". *Twelfth Workshop on Standards for the Interoperability of Distributed Simulations*, pp. 1995
- [13] Marshall, D., A. McCoy, D. Delaney, S. McLoone and T. Ward "A Realistic Distributed Interactive Application Testbed for Static and Dynamic Entity State Data Acquisition". *IEEE Irish Signals and Systems Conference*, Belfast, Ireland, pp. 83-88., 2004