Analysing movement and world transitions in virtual reality tele-conferencing

Chris Greenhalgh
Department of Computer Science, University of Nottingham, U.K.
c.greenhalgh@cs.nott.ac.uk

In this paper we make use of automatically generated logs of user activity from 6 meetings held using the MASSIVE-1 virtual reality tele-conferencing system to determine a number of characteristics of user movement and world transition. These results are applied to a consideration of four issues for CVE system design and resource requirements: the amount of network bandwidth and computation required to handle movement within worlds; the degree of look-ahead required when moving; whether world transitions by groups of participants could benefit from special handling (e.g. some form of multicast state-transfer); and whether caching of world state would be useful in these contexts. In each case the implications are quantified for the meetings analysed.

1. Introduction

Within the BT/JISC-funded Inhabiting The Web (ITW) project 17 meetings have been held using the MASSIVE-1 virtual reality tele-conferencing system. MASSIVE-1, which was presented at ECSCW'95 (Greenhalgh and Benford, 1995a), supports real-time 3D graphical, audio and text interaction between geographically isolated users at graphical workstations and is based on the spatial model of interaction, which was presented at ECSCW'93 (Benford and Fahlen, 1993). In MASSIVE-1, each participant is (graphically) embodied and can move freely and independently within a connected universe of virtual worlds, talking to other participants that they meet.

MASSIVE-1 has been the subject of at least three previous analyses: our own reflection on the use of the system (at ECSCW'95); and two studies by Bowers, O'Brien and Pycock (1996a and 1996b) which employ social scientific methods (employing talk and virtual body movement transcription) and ethnographics ob-

J Hughes et al. (eds), Proceedings of the Fifth European Conference on Computer Supported Cooperative Work, 313-328.

servation to investigate social interaction in CVEs, and the concurrent interactions which occur in the real and virtual worlds. All of these analyses have been qualitative, and have focused on particular (characteristic) incidents or experiences.

In contrast, in this paper we make use of automatically generated logs of user activity in the virtual worlds, which are analysed mathematically to explore the full range and occurrence of movement and world-transition behaviours over a relatively long period of time (6 meetings). This analysis of overall characteristics has direct application to system design and network requirements for CVEs in general. In this paper we show how this information can help to determine:

- the amount of network bandwidth and computation required to handle movement within worlds;
- the degree of look-ahead required when moving;
- whether world transitions by groups of participants could benefit from special handling (multicast state-transfer), and what delay this would introduce; and
- whether caching of world state would be useful in these contexts and on what time-scale.

In the next section we give some further information about the meetings which occurred, and which the subsequent analysis is based on. Section 3 describes the nature and handling of the data. Section 4 presents the direct results of the data analysis, while section 5 goes on to consider the four questions above. Finally, section 6 presents the overall conclusions and acknowledgements.

2. The meetings

This section gives some details of the nature and form of the meetings which occurred over the course of the ITW project using the MASSIVE-1 CVE.

Over the course of the project 17 official virtual meetings were held using MASSIVE-1 over SuperJANET, the UK's high-speed academic network. The meetings involved between 5 and 9 simultaneous participants based at BT and at (some or all of) the five participating UK Universities: Nottingham, Lancaster, Manchester, Leeds and UCL.

Each trial lasted approximately one hour. The first 11 were concerned primarily with familiarising the participants with the technology and with ensuring consistent and acceptable performance of the system. The remaining 6 meetings combined project management activity (for the ITW project itself), informal social interaction and organised games and activities. The data in this paper is based on these last 6 meetings. For the 6 meetings in question, in addition to project management, the organised activities and meeting attendance were:

- team word games (14th August 1996, 8 participants, 6 sites);
- team anagrams and other word games followed by hide and seek in the maze (21st August 1996, 9 participants, 6 sites);
- multi-world exploration and problem solving in teams followed by a maze

race (4th September 1996, 5 participants, 4 sites);

- team exploration of a noisy environment requiring control of focus (18th September 1996, 6 participants, 4 sites);
- two individual games (balloon debate and discussion) (25th September 1996, 8 participants, 5 sites); and
- an end-of-trials party with a dance contest, obstacle race and discussion (2nd October 1996, 8 participants, 6 sites).

In the next section we consider in more detail the actual information captured during meetings and way in which it has been analysed.

3. The data

Having described something of the nature of the meetings that were held, we now describe the information that has been captured and the way in which it has been used in this paper. There are five main sources of data about the meetings:

- log files generated by the MASSIVE-1 user client programs;
- network logs for audio data captured using the UNIX utility "tcpdump";
- videos of the meetings from the perspective of one or more participants;
- questionnaires completed after the meetings; and
- personal reflections on the meetings from those involved.

This paper is based primarily on the first of these only, since they are amenable to automated analysis and include all of the information required for this analysis of the movement and world transition. Specifically, these log files from the MASSIVE clients include information about: position and orientation within worlds; and transitions between worlds.

Log files are available for approximately 70% of participants at the 6 meetings in question. For the 6 meetings there are a total of 35.9 hours of log files (from individual participants) covering approximately 8 hours of meetings.

3.1. Example meeting: 25th September 1996

Figure 1 illustrates the available log data for the meeting held on 25th September 1996. Time runs from left to right with time of day values recorded at the bottom of the graph. At the left is a column of labels which identify virtual worlds ("presentation-world", "balloon", etc.), and within each virtual world identify participants (by the IP number of the machine which they were using, e.g. "128.16.8.225", "128.243.22.16"). Thus the horizontal area at the top of graph between the first pair of horizontal lines represents activity in the "presentation-world"; in this case the participant using machine "130.88.13.171" enters that world at about 15:10 from "meeting-world" and returns almost immediately. The second horizontal area represents activity in the world "balloon", and so on. The activity of each participant is recorded by a horizontal trace which evolves with time from left to right. When

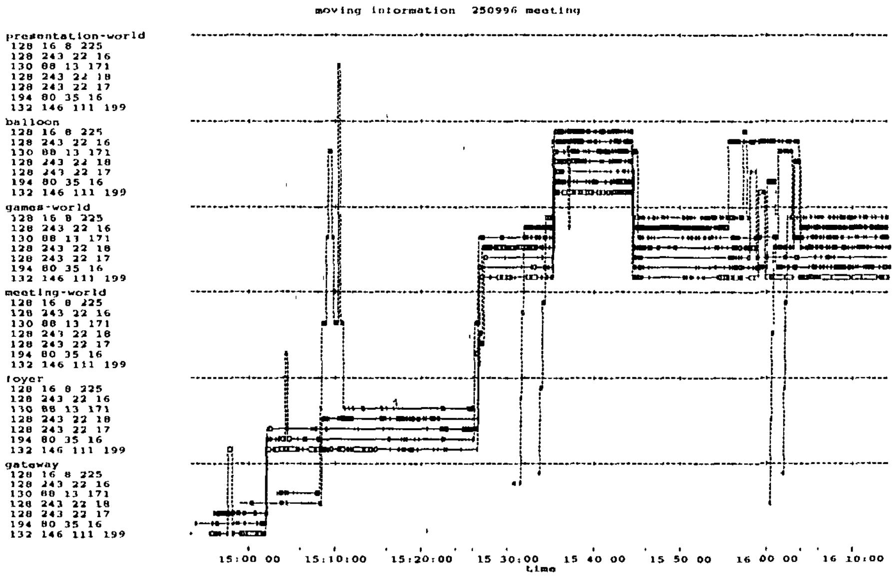

Figure 1. Visualisation of data from 25/9/96 meeting.

the person is stationary the trace is a thin line; when the person is moving the trace becomes a thick line or a box; when the person moves directly to another world (via a portal) this is indicated by a dashed vertical line.

Key events which may be seen from the graph include: participants arrive in "gateway" from 14:53 to about 15:02, when most of them move into the "foyer" world; the project management phase of meeting runs from then to about 15:26, when everyone moves on to the "games-world"; the balloon debate runs from about 15:35 until 15:44 in the "balloon" world and is followed by a discussion and the introduction of next task which runs from about 15:55 to about 16:04; the meeting ends at about 16:14 when everyone finally leaves the world (and the system).

3.2. Data handling

All data processing and analysis was done using in-house tools developed for the purpose. The key events used in the analysis are: participants entering and leaving the system; participants moving between worlds; and participants starting and stopping moving. All of the results are presented as distributions of measures extracted from the data, such as time spent moving, time spent in a world, etc.

The next section presents the actual measures derived from the data, i.e. the basic results, while section 5 goes on explore the implications of these results.

4. Results

In this section we present the basic results of this paper, based on the analysis of logged data from 6 meetings. The specific issues addressed by the analysed data are:

- How much of the time do participants spend moving?
- Do people all move simultaneously or do they move independently?
- How quickly do participants move about within the virtual worlds?
- Do participants move to new worlds in groups or individually?
- If moving between worlds in groups, how spread out are those groups?
- Do participants revisit worlds, and if so with what time lapse?

In the section 5 we go on to consider the implications of these results for CVE system design and resource requirements. While these results form the foundation for the implications it is possible to read skip directly to section 5 and refer back to individual results as they are considered in that section. In this section we take each of the above issues in turn.

4.1. How much of the time do participants spend moving?

The first results concern the fraction of the time which individual participants spend moving while in the virtual world. I.e. how long do people spend moving and how long are they stationary? MASSIVE-1 uses a mouse-based navigational interface in which participants click on different parts of the graphical window to move in different directions.

For all participants and all six later meetings (subject to available log files) the average percentage of time spent moving was 19.6%. Figure 2 shows the distribution of percentage of time spent moving for a single participant during a single visit to a single world. I.e. each visit of a participant to a world has been considered independently, and a percentage value calculated for the fraction of time which they spent moving while in that world on that occasion. The upper curve is the cumulative distribution while the lower jagged line is the point distribution with a bucket size of 1%.

Analysing the same data by participant over all six meetings (all worlds visited) reveals that average movement rates for individual participants vary from 7.2% to 28%. Similarly, average movement rates for particular worlds (all participants) vary from 7.5% to 54.6%. Both of these extremes occurred in the final meeting which included the end-of-project virtual party; the lowest movement rate occurred in the classical music world while the highest movement rate occurred in the disco world (home of the disco-dancing competition). The average movement rate in the formal meeting world was 16.7%.

This result, like the others, is simply presented in this section: reflection is reserved until section 5.

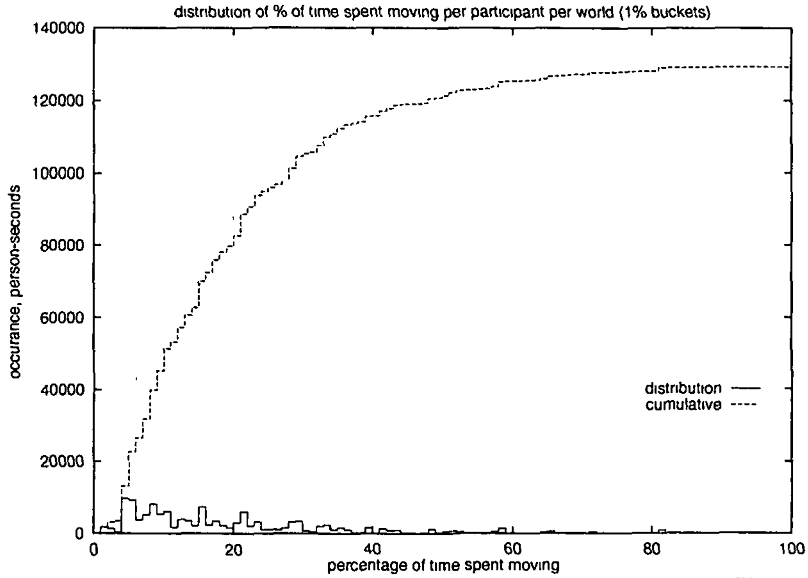

Figure 2. Distribution of percentage of time spent moving, all log files.

4.2. Do people all move simultaneously or do they move independently?

The second result considers whether people in a single world tend to move simultaneously or not. I.e. if one person is moving is everyone likely to be moving? (following?)? Or are people more likely to be stationary instead? (watching?) Or is there no connection?

For all six meetings and all participants the groups or participants simultaneously present in a world has been considered to establish the distribution of numbers of participants moving concurrently. Figure 3 shows this distribution (the solid line). There were over 9000 seconds of logged activity in which exactly one person was moving in a world. The incidence of larger numbers of participants moving simultaneously falls off smoothly. For comparison, the dashed line shows the distribution that would be expected if movement were completely independent, e.g. random.

It is apparent that the incidence of larger numbers of people moving concurrently is rather higher than chance, indicating that (at least some of the time) people coordinated their movement activities. This would clearly be the case when, for example, the meeting moves together from one world (and activity) to another.

4.3. How quickly do participants move about within the virtual worlds?

The third result concerns the speed with which participants move about in the virtual environment. Movement in MASSIVE-1 is not constrained by any particular

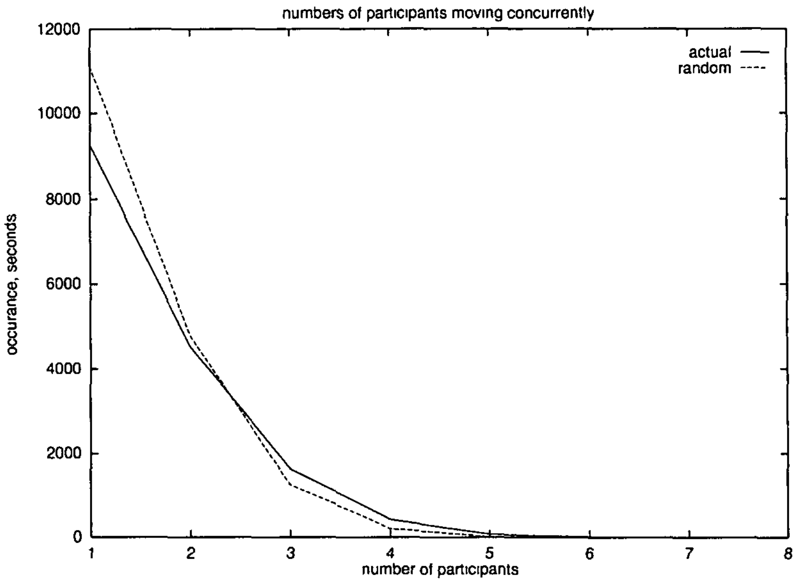

Figure 3. Distribution of numbers of people moving concurrently.

physical model or real-world analogy. According to where on the screen users click the mouse buttons they move at different speeds. The nominal maximum speed is 8 units per second, though higher speeds may occur as a result of uneven frame rates (i.e. where drawing subsequent graphical images takes significantly different lengths of time). A unit is approximately equivalent to a metre (based on scale size of embodiments and world geometries).

Figure 4 shows the cumulative distribution of speeds of movement for all participants in all (6) meetings. The average speed overall (including periods when stationary) is 0.18 units per second. The average speed when moving is 1.04 units per second. The distribution is a smooth curve with higher speeds decreasing in likelihood. Note that the vertical axis of the graph does not start at zero: the high incidence of being stationary (80.4% of the time - see section 4.1) is not displayed.

4.4. Do participants move to new worlds in groups or individually?

The fourth measure concerns the question of whether participants tend to move between worlds in groups or as isolated individuals. From the nature of the meetings we expect that group transitions will occur, as when the meeting moves from the "games-world" to "balloon" world in the example meeting (section 3.1). Here we will quantify this effect to help us establish its general significance (the next result will address the spread of group transitions).

For this measurement we have identified the occurrences of two or more participants moving from one world to another so that at some point in time all are in one world and at a later point in time all have moved directly to the other world and are

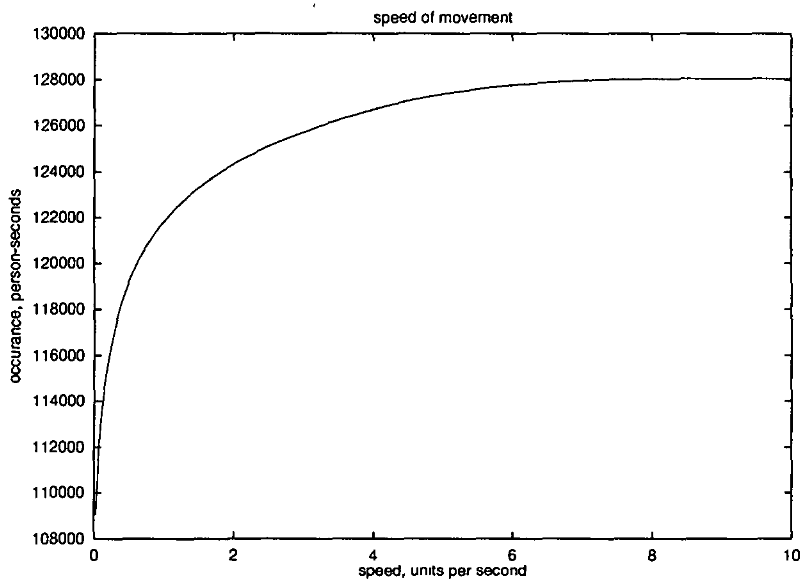

Figure 4. Distribution of speed of movement.

still there. This reflects a group moving together to another world. Figure 5 shows the incidence of group world transitions for different group sizes. The upper

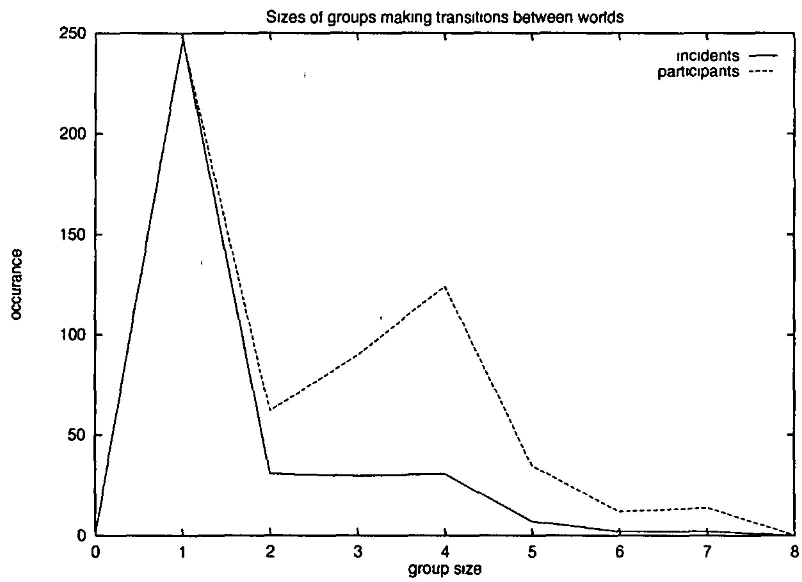

Figure 5. Distribution of sizes of group moving between worlds.

(dashed) line is weighted by the number of participants in the group, while the lower line shows just the number of incidents. The upper curve is representative of the importance of this effect as it the number of participants involved A participant entered a new world 584 times in the available data. 337 (58%) of these person-transitions were in groups of 2 or more participants. In total there were 350 incidents: 247 (71%) solitary transitions and 103 (29%) group transitions occurred. The average group size for group transitions was 3.27; for all transitions (including solitary transitions) the average group size was 1.67. The largest group making a transition had 7 members (constrained by available log files).

4.5. If moving between worlds in groups, how spread out are those groups?

Having established above that indentifiable groups move between worlds, we now consider how spread out in time those groups are. I.e. from start to finish, how long does it take a group to make a world-transition through a single portal. There is no collision detection in MASSIVE-1, so that participants could in principle all move through the portal simultaneously.

Figure 6 shows the distribution of world entry delays for group members lagging behind the first group member to make the transition. I.e. the time on the graph is

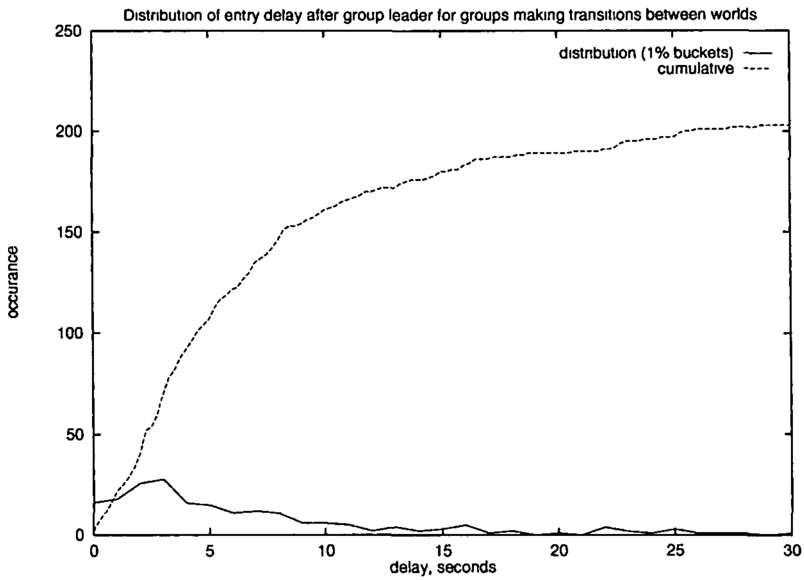

Figure 6. Distribution of world entry delay for group members.

the time which has elapsed since the first member of the group entered the new world. The graph shows delays of up to 30 seconds which account for 203 (87%) of the 234 participants who follow another participant into a world. 104 (44%) of these occur within 5 seconds while 159 (67%) occur within 10 seconds of the first group member making the transition.

4.6. Do participants revisit worlds, and if so with what time lapse?

The last result concerns participants who visit a world more than once during a single meeting. In particular, when a participant does visit a single world twice in a meeting, how much time elapses between those visits?

Figure 7 shows the cumulative distribution of times between repeat visits to a world for all participants in all meetings. Return visits occurring after more than 24

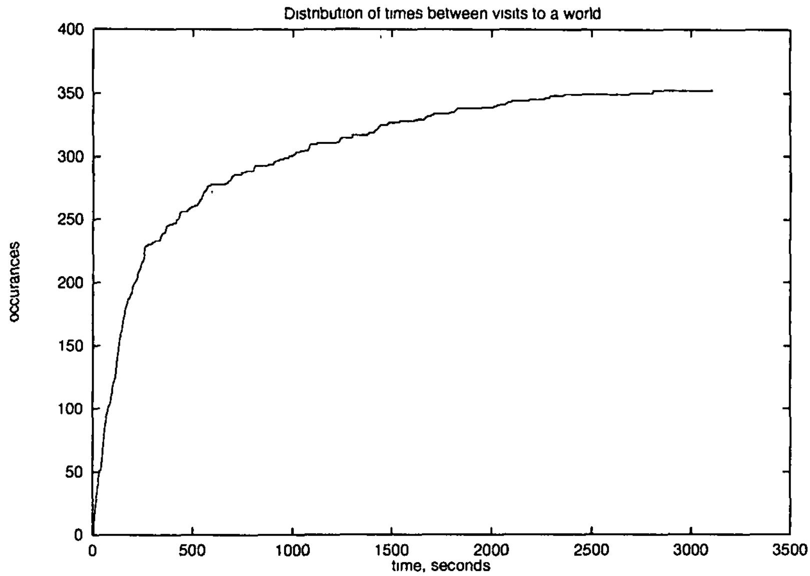

Figure 7. Distribution of times between return visits to a world.

hours are not shown (i.e. only return visits occurring in the course of a single meeting are shown). There are 353 occurrences in total, with an average delay of 427 seconds (just over 7 minutes) between visits. Half of these return visits occur within 162 seconds (less than 3 minutes). 231 visits to worlds do not feature in this distribution because they were the first visit to a world during that meeting, i.e. 60% of world transitions were return visits within a single meeting.

We have now seen all of the basic results obtained from analysing the logged statistics from 6 of the ITW meetings. We now move on to consider the implications and applications of these results.

5. Implications

In this section we use the results from the previous section's analysis of the meeting data to reflect on four implementation issues for CVEs: the amount of network bandwidth and computation required to handle movement within worlds; the degree of look-ahead required when moving; whether world transitions by groups of participants could benefit from special handling; and whether caching of world state would be useful in these contexts. We consider these issues in turn, and refer to individual results from section 4 as appropriate.

5.1. Handling movement within worlds

The first issue which we consider is that of handling movement with virtual worlds: handling movement in virtual worlds consumes network and computational resources since user movements are non-deterministic and so must generate network traffic to inform other participants, and this traffic and consequent world cache updates require computation resources from recipients. Consequently, when assessing the number of participants which a system may support or the resources required to support different numbers of participants it is important to know how much participants move and whether the system should be expected to cope with all participants moving simultaneously.

From section 4.1 we observe that in the meetings considered participants move on average 19.6% of the time, and we might use this figure in calculating expected average bandwidths produced by a large number of participants. (The model of network traffic for MASSIVE-1 described by Greenhalgh and Benford (1995b) used a value of 25% based on earlier in-lab use MASSIVE-1 for lab meetings.) This aspect of movement traffic behaviour is normally ignored when modelling resource requirements of CVEs, and a single long-term average traffic rate used, (e.g. Macedonia, Zyda, Pratt, Barham, and Zeswitz,1994).

From section 4.2 we observe that participants do move together more often than chance within common-purpose groups. For example, 6 participants moving simultaneously was observed for 11.3 seconds in the test data compared with an expected value of 1.1 seconds for independent movement. So the system should be able to cope with simultaneous movement of members of coherent groups being more common that might be expected for completely independent (and isolated?) participants.

The above values will be affected by techniques such as update rate limiting and dead reckoning which reduce the (worst case) amount of traffic needed to communicate movement information. For example, if positional updates are sent no more than once per second then the effective time spent moving is increased to 33.4% (from 19.6%). This is because short periods without movement are effectively swallowed up by the surrounding periods of movement.

5.2. Spatial look-ahead

The second issue with design implications is spatial look-ahead in virtual worlds: to enhance the scalability of large virtual worlds some form of spatial coherence can used (e.g. NPSNET (Macedonia, Zyda, Pratt, Brutzman and Barham, 1995), MASSIVE-2 (Benford, Greenhalgh and Lloyd, 1997), PARADISE (Singhal and

Cheriton, 1996)). Depending on the organisation of the system more distant participants and parts of the environment may need to be acquired (e.g. paged into memory) as a participant moves about. This paging may take a significant period of time (e.g. several seconds) especially when it must be acquired across the network or it contains a very large amount of information. Therefore to avoid or reduce abrupt discontinuities in the participants view of the environment it may be necessary to anticipate their future requirements and to begin paging in regions before they are needed. The degree of anticipation required will depend (amongst other factors) on the speed with the participant is expected to move.

In the meeting considered here, which involved relatively small virtual spaces comparable in size to large rooms or small halls, the average rate of movement was 1 virtual metre per second (section 4.3) - a slow walking pace. Participants had interfaces which allowed them much greater speeds (at the expense of accuracy), and speeds of 6-8 (the nominal maximum) virtual metres per second are uncommon (occurring only 0.25% of the time). We may suggest from this that in normally scaled room structures people adopt a normal speed of movement - approximately walking pace, and look-ahead may be tailored accordingly.

However it is unlikely that this result will generalise to large or exterior spaces. After all, the average speed of movement on a road is more like 25 metres per second (and this often feels frustratingly slow).

5.3. Group world transitions

The third issue we consider concerns group world transitions and the potential utility of multicast state transfers.

Virtual reality systems often exploit a metaphor of multiple disjoint virtual worlds connected by gateways or portals (e.g. DIVE (Carlsson and Hagsand, 1993), MASSIVE-1 and VRML worlds). On moving to a new world a process (e.g. a participants interface) must acquire initial information about that new world. This may be by means of a unicast state transfer from a peer process (as in DIVE) or processes (as in MASSIVE-1); it may be by means of a unicast state transfer from a server (as in VRML); it may emerge over time from watching activity in the new world (as in DIS 2.x with periodic entity state transmissions). If groups of participants tend to move to other worlds together than there is a clear opportunity to make use of multicasting where a specific state transfer phase is present. This is especially true if those groups can be identified in advance, for example within another virtual world heading towards a portal.

In the six meetings under consideration there were 584 occurrences of a participant entering a new world, and of these 337 (57%) involved travelling in a group of 2 or more (section 4.4). It can be argued that the fact that group transitions occur is just a reflection of the application and the way in which the meetings were organised - a different application of CVEs might not involve (coordinated) world transitions at all. On the other hand, it can be argued that structuring the meeting into distinct phases involving different virtual worlds was found to be useful and effective and so is likely to be used in a broad range of applications. Also, group transitions may be observed within structured worlds as well as between worlds (e.g. entering a room or building, leaving a concert - see (Benford, Greenhalgh and Lloyd, 1997)) and the same techniques will be relevant.

Given that group transitions occur (in whatever context) it is also important to understand how well defined they are and how long they take. For example, if a multicast state transfer were employed then all group members would effectively reach the new world (or region) at the same time, although they may have set off at different times. This will necessarily delay the earlier members and/or require speculative state transfers to those participants who are *expected* to follow. Referring to section 4.5, if every group transition used multicast state transfer then 337 unicast state transfers would be replaced by 103 multicast state transfers (eliminated 69% of unicast state transfers), however the longest resulting delay would be 13 minutes! More realistically, we might constrain the maximum delay (i.e. the period over which group members may arrive and still use the same multicast state transfer). For delays up to 5 seconds 104 unicast state transfers would be eliminated (30.8%); for delays up to 10 seconds this rises to 159 (47%); for delays up to 30 seconds it rises to 203 (60%). No longer delay is likely to be acceptable except in very specific applications.

On the negative side, the management overhead for the multicast state transfer may well include a component for every recipient; so multicast state transfers will be increasingly attractive where larger quantities of state have to be transferred. MASSIVE-2 uses multicast state transfers when participants arrive in quick succession, and unicast transfers otherwise. However performance data is not yet available to assess the effectiveness of this implementation.

5.4. World state caching

The final issue which we reflect on concerns the possibility of caching world state between visits. CVE systems which use state transfer on world (or region) may repeat this transfer each time a participant returns to the same world (or region) (e.g. MASSIVE and DIVE both do this). Clearly, the participant could (resources permitting) retain a local cache of some or all of this state information between visits, alleviating or eliminating the requirement for a new state transfer. In order for this to be effective some or all of the state information from the previous visit must still be valid. I.e. the world cannot have changed; or if it has changed then it must have changed only in limited and well-defined ways which can be updated independently. Shortage of resources (for caching) and the potential volatility of world information may dictate that caching be limited in extent and duration. Information from the ITW meetings is relevant to this second limitation of duration. Ideally, we would like to keep information only for those worlds which will be returned to. One heuristic by which we can attempt to consider the likelihood of returning to a world is the time since it was last visited (this will also impact the likely usefulness of the cached data).

In the meeting data under consideration there were 353 occurrences of participants returning to a previously visited world within the course of a meeting (section 4.6). The average delay was just over 7 minutes, but half of the incidents occurred within 3 minutes of leaving the world. This discrepancy between mean and median indicates that returning to a world within a few minutes is relatively common, but that above about 6 minutes (67% of occurrences) the return times spread out much more. I.e. we can make a rough division between return visits which happen within a few minutes and which account for the majority of incidents, and the remainder which happen over much longer periods. In the appropriate circumstances this could be an appropriate cut-off point for short-term caching of worlds visited.

Again, it can be argued that we are just seeing the effects of the particular choice of meeting agenda and universe structure. Certainly, more data would be needed to generalise these results to a broader range of scenarios. On the other hand at least some of the return visits reflect the choice to structure the meeting worlds as common hub or access worlds with multiple activity worlds "hanging" off them. This is a general structuring technique (c.f. hierarchical file stores) which might be expected to recur in many systems and applications, and to encourage (or require) repeated return to the hub world(s), making this effect wide-spread.

This completes our reflection on the implications of the participant movement and world-transition data obtained from the later ITW virtual meetings. The final sections presents our conclusions in a more succinct form and outlines future work.

6. Conclusions and future work

In this paper we have analysed data captured automatically during the BT/JISC-funded Inhabiting The Web (ITW) trials of the MASSIVE-1 virtual reality tele-conferencing system. These concern structured meetings with 5-9 participants in a combination of group management, unconstrained social interaction and organised games and activities. Using this data we have produced answers to the questions:

- How much of the time do participants spend moving?
  About 20%, though there is considerable variation between individuals (7% to 28%) and between worlds (7% to 55%). This value is also sensitive to bandwidth reduction techniques like rate-limiting and dead-reckoning.
- Do people all move simultaneously or do they move independently?
  People move simultaneously more than we would expect from independent movement, though the drop off in incidence of simultaneous movement with numbers is still very marked.
- How quickly do participants move about within the virtual worlds?
  About walking pace (1 virtual metre per second on average) within the room-scaled environments used in the ITW project.

- Do participants move to new worlds in groups or individually?
  They move in groups on average 58% of the time, the average group size being 1.67 (or excluding solitary transitions, 3.27). The largest size group observed making a world transition was 7 participants.
- If moving between worlds in groups, how spread out are those groups?
  Transition times of up to 5 seconds account for 44% of the participants who follow another participant into a world; times up to 10 seconds account for 67%; and time up to 30 seconds account for 87%.
- Do participants revisit worlds, and if so with what time lapse?
  60% of participant visits to worlds were return visits during the course of single meeting. The average return time was just over 7 minutes, though 50% of return visits occurred within 3 minutes and 67% of return visits occurred within 6 minutes.

The four design and resource issues which we have considered are:

- Network bandwidth and computational requirements to handle movement within worlds -
  must account for movement at least 20% of the time (depending on the use of other bandwidth-limiting techniques and the individual character of the world), with an independent movement model being unsuitable for coherent groups of individuals.
- The degree of look-ahead required when moving -
  may use normal "real-world" movement speeds as a heuristic expectation in small interior virtual worlds.
- The scope for using multicast to handle transitions by groups of participants -
  there is significant scope for using multicast state-transfers to handle group world transitions which accounted for 58% of world transitions; delays in the range 5-30 seconds may be chosen, trading off delay and effectiveness of multicasting (44%-87%).
- Whether caching of world state would be useful in these contexts and on what time-scale -
  caching of world state would have accounted for 60% of world transitions in these meetings, and a relatively short cache expiry time of 6 minutes would have accounted for the majority of these (67%); this is clearly potentially beneficial.

The scope of these answers will be limited by the specific range of activities occurring in the meetings under consideration. Future work is needed to extend these results and reflection to include:

- larger numbers of simultaneous participants;
- other uses or applications, e.g. spontaneous meetings and unstructured (or self-directed) use; and
- other virtual geographies and topologies such as large exterior spaces.

Other issues which remain to be addressed from this same data set may include:

- characteristics and requirements of audio interaction;

- suitability of dead-reckoning and other movement bandwidth limiting techniques for user movement without simulated "real-world" constraints; and
- characterisation of network traffic and behaviour for further modelling, including consideration of burstiness and systematic variations (e.g. with task).

Acknowledgements

We would like to gratefully acknowledge the support of British Telecommunications plc. in funding and participating in this research, which has been carried out within the BT/JISC funded Inhabiting The Web (ITW) project. With thanks to all those who participated in the meetings Jolanda Tromp, Gail Reynard, Steve Benford, Rob Ingram, David Lloyd and Lucy Smallwood at the University of Nottingham; Dave Small at the University of Leeds; John Cook and Adrian West at the University of Manchester; Angeline Lim at UCL; Jason Morphett, Paul Shumake, Charanjit Sidhu and Natasha Upal at BT, and especially Adrian Bullock at Nottingham for making sure that the meetings happened and for collecting the data.

References

Benford, S. and Fahlen, L. (1993): "A Spatial Model of Interaction in Virtual Environments", in *Proc. Third European Conference on Computer Supported Cooperative Work* (ECSCW'93), Milano, Italy, September 1993.

Benford, S., Greenhalgh, C., and Lloyd, D. (1997): "Crowded Collaborative Virtual Environments", in *Proc 1997 ACM Conference on Human Factors in Computer Systems* (CHI'97), Atlanta, Georgia, March 22-27 1997, ACM Press.

Bowers, J., Pycock, J. and O'Brien, J. (1996a): "Talk and Embodiment in Collaborative Virtual Environments", in Proc. CHI'96, New York: ACM Press, 1996.

Bowers, J., O'Brien, J. and Pycock, J. (1996b) "Practically Accomplishing Immersion: Cooperation in and for Virtual Environments", in *Proc. CSCW'96*, New York. ACM Press.

Carlsson, C. and Hagsand, O. (1993): "DIVE - A Multi-User Virtual Reality System", in *Proc. VRAIS'93, IEEE Virtual Reality Annual International Symposium*, pp. 394-400.

Greenhalgh, C., and Benford, S. (1995a): "Virtual Reality Tele-conferencing: Implementation and Experience", *Proc. Fourth European Conference on Computer Supported Cooperative Work* (ECSCW'95), Stockholm, September, 10-14 September, 1995, Kluwer Academic Publishers, Dordrecht, pp. 165-180.

Greenhalgh, C. and Benford, S. (1995b): "MASSIVE: A Virtual Reality System for Tele-conferencing", *ACM Transactions on Computer Human Interfaces* (TOCHI), Volume 2, Number 3, pp. 239-261, ISSN 1073-0516, ACM Press, September 1995.

Macedonia, Michael R., Zyda, Michael J., Pratt, Donald P., Barham, Paul T., and Zeswitz, Steven (1994): "NPSNET: A Network Software Architecture for Large-Scale Virtual Environments", in *Presence*, Vol. 3, No. 4, Fall 1994, pp. 265-287, MIT Press.

Macedonia, Michael R., Zyda, Michael J., Pratt, David R., Brutzman, Donald P., and Barham, Paul T. (1995): "Exploiting Reality with Multicast Groups: A Network Architecture for Large-scale Virtual Environments", in *Proc. 1995 IEEE VRAIS'95*, 11-15 March, 1995, RTP, North Carolina.

Singhal, Sandeep K., and Cheriton, David R. (1996): "Using projection aggregations to support scalability in distributed simulation", in *Proc. of the 1996 International Conference on Distributed Computing Systems*, IEEE.