2006

Minimising the computational cost of providing a mobile immersive communication environment (MICE)

Ying Peng Que
University of Wollongong, ypq01@uow.edu.au

Paul Boustead
University of Wollongong, boustead@uow.edu.au

Farzad Safaei
University of Wollongong, farzad@uow.edu.au

Minimising the computational cost of providing a mobile immersive communication environment (MICE)

Abstract

This paper investigates the minimisation of the computational cost of providing Immersive Voice Communication (IVC) for mobile clients attached to a Distributed Virtual Environment (DVE). When experiencing IVC, each listening avatar in the DVE, receives a mix of the surrounding avatars' voices, all rendered according to their respective positions in the virtual world. We propose to deliver our mobile IVC service using a type of server-based network architecture to perform voice spatialisation function and mixing for each mobile client. This model will minimise both the access bandwidth and processing power requirements of the mobile devices. We propose a model for computation reduction in servers that re-uses the result of audio spatialisation for several clients by allowing certain level of voice localisation error. Using Linear Programming (LP), we have developed two optimisation formulations to examine different scenarios of applying the computation reduction scheme, with different avatar densities and under varying rendering accuracy constraints. The insights gained from our analytical results should benefit future IVC service providers for mobile DVE.

Disciplines

Physical Sciences and Mathematics

Publication Details

This paper originally appeared as: Que, YP, Boustead, P & Safaei, F, Minimising the computational cost of providing a mobile immersive communication environment (MICE), CCNC 2006. 3rd IEEE Consumer Communications and Networking Conference, 8-10 January 2006, vol 2, 1163-1167. Copyright IEEE 2006.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE CCNC 2006 proceedings.

Minimising the Computational Cost of Providing a Mobile Immersive Communication Environment (MICE)

Ying Peng Que, Paul Boustead, Farzad Safaei
Smart Internet CRC,
Telecommunications and Information Technology Research Institute,
University of Wollongong, Australia
Email {ying, paul, farzad}@titr.uow.edu.au

***Abstract*-This paper investigates the minimisation of the computational cost of providing Immersive Voice Communication (IVC) for mobile clients attached to a Distributed Virtual Environment (DVE). When experiencing IVC, each listening avatar in the DVE, receives a mix of the surrounding avatars' voices, all rendered according to their respective positions in the virtual world. We propose to deliver our mobile IVC service using a type of server-based network architecture to perform voice spatialisation function and mixing for each mobile client. This model will minimise both the access bandwidth and processing power requirements of the mobile devices. We propose a model for computation reduction in servers that re-uses the result of audio spatialisation for several clients by allowing certain level of voice localisation error. Using Linear Programming (LP), we have developed two optimisation formulations to examine different scenarios of applying the computation reduction scheme, with different avatar densities and under varying rendering accuracy constraints. The insights gained from our analytical results should benefit future IVC service providers for mobile DVE.**

Keywords: Virtual Environments for Entertainment, Wireless and Mobile Gaming, Mobile and Wireless Entertainment

I. INTRODUCTION

In recent time, the delivery of networked voice communication services in Distributed Virtual Environment (DVE) has attracted much research interest [1] [2] [3]. DVE is also known as Networked Virtual Environment (NVE) or Collaborative Virtual Environments (CVE). A DVE is a virtual environment which is distributed over a common underlying network [1]. Each DVE user is represented by an avatar in the virtual world. One typical example of DVE is Multi-player Online Games (MOG) such as Lineage II which had 2.1 million subscribers in January, 2005 [4]. In a DVE, multiple users concurrently explore the virtual world which is often represented as high quality visual scenes on the users' computers. However, close interactions and co-operations among avatars are likely to require multi-party voice communication. Hence a multipoint-to-multipoint Immersive Voice Communication (IVC) system could be considered very useful for a DVE. The IVC system aims to create a personalised *auditory scene* for each listening avatar comprised of all the speaking avatar voices within its hearing range, each rendered with cues for their respective direction and distance. In many cases, a DVE needs to cater for the concurrent accesses of a large number of avatars and the distribution of avatars can be quite dense. More importantly, the DVE avatars can be very close in the virtual space but yet spread over a large geographical scale in the physical world. It is therefore important for IVC system to achieve a good balance between its rendering quality and scalability. Two important constraints that a scalable IVC system must satisfy are the limited access bandwidth of the client devices and the scarce local computational resources on these devices.

One special class of DVE is mobile DVE in which users access the virtual environment using handheld wireless devices such as mobile gaming platforms exemplified by the SONY PSP and Nintendo DS systems. We refer to our IVC service in mobile DVEs as a Mobile Immersive Communication Environment (MICE). For mobile devices, the two aforementioned scalability constraints become more stringent. This is due to the challenges in incorporating high speed processing technologies into compact mobile platforms and the scarce transmission bandwidth in wireless networks. Furthermore, mobile platforms are constrained by their battery power supplies which limit the duration and complexity of the on-board processing operations. Finally, the user access platforms in MICE often require portable voice playback mechanism such as headsets. In order to minimise the computational load on the mobile client platforms in MICE, we propose a server-based architecture in which, *auditory scene creations* are carried out centrally at the servers within the fixed infrastructure by the service provider. We propose to use the *Head Related Transfer Function (HRTF) localisation* technique. The term localisation here refers to defining the direction of avatar voices in 3-D space. *HRTF localisation* uses portable binaural headphones for final voice playback [5]. Our architecture offers much better access bandwidth efficiency than the peer-to-peer architecture reviewed in [1]. Each user in our architecture sends a single mono stream to the corresponding server. Due to the use of *HRTF localisation*, servers in our architecture only need to send two mixed rendered streams (left and right channels [5]) to each client for final playback. In our current work, we have simulated MICE over a centralised server architecture as described in [1].

II. OVERVIEW OF MICE

A. Auditory Scene Creation for MICE

The *auditory scene creation* process for MICE consists of two stages. In the first stage, the voice of each speaking avatar is localised (having its direction defined by *HRTF localisation*) for the respective virtual world positions of all the listening avatars within the corresponding audible range. In the second stage of *auditory scene creation*, the rendered streams of all the speaking avatars in the hearing range of each listening avatar are distance-weighted and mixed together for final playback. Due to the aforementioned computational power constraint, the distance-weighting operation uses simple

amplitude weighting according to the inverse square law of sound propagation through free-space. However, sophisticated models of sound reflections and reverberations [6] can be incorporated into our IVC system if processing resources are available. In the mixing operation, the distance-weighted and rendered streams are linearly mixed. All the left channel streams are mixed into one final left channel stream and similarly all the right channel streams into one final right channel stream. Fig. 1 illustrates the creation of the auditory *scene* for listening avatar L_1 through mixing the distance-weighted rendered streams of four speaking avatar voices (S_1-S_4).

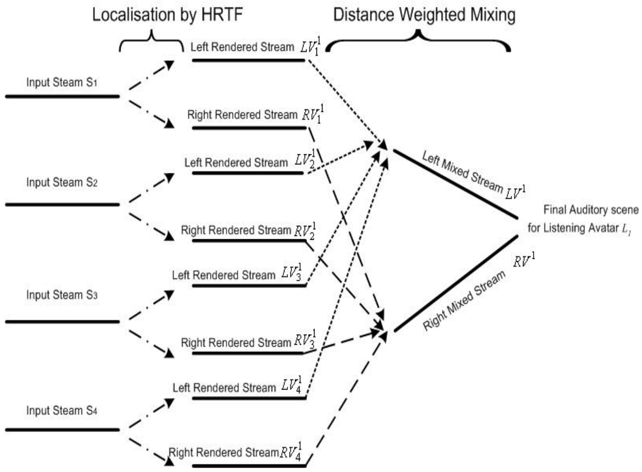

Fig. 1 The "brute force" approach to *auditory scene creation*

B. Distance-governed Computation Reduction

Fig. 1 illustrates the *"brute force"* approach of creating MICE where the voice of each speaking avatar is rendered by *HRTF localisation* according to the exact positions of all the surrounding listening avatars. Due to the high computational cost of *HRTF localisation* [5], such *"brute force"* approach of creating MICE can be prohibitively expensive especially when the avatar distribution is dense. It is hence important to develop effective techniques to reduce the computational cost incurred by voice rendering using *HRTF localisation.* The distance-governed computation reduction scheme algorithms proposed here seek to minimise the number of HRTF localisations required for a particular MICE.

Let S_i denote any given speaking avatar where $1 \le i \le B$ (B denotes the total number of speaking avatars in a MICE). Let L_j denote any given listening avatar where $1 \le j \le A$ (A denotes the total number of listening avatars in a MICE). We define avatar density as the average number of speaking avatars in the hearing range of a given listening avatar. In our work, we use the term *directional vector* to refer to the rendered voice of a particular speaking avatar along a particular direction with the position of the speaking avatar at the origin (0, 0). Let N denote the maximum number of possible *directional vectors* around a particular speaking avatar S_i. We initially set N to cover all 360 degrees on a 2-D plane around a particular speaking avatar. However N can be decreased if a coarser granularity in localisation is acceptable. We use v_i^k to denote a particular *directional vector* for the i^{th} speaking avatar in direction k, where $1 \le k \le N$. For example, the *directional vector* v_1^{220} shown in Fig. 2, is the voice of the speaking avatar S_1 rendered along 220 degrees, with S_1 at the (0,0).

The number of HRTF localisations within a MICE, equals to the number of directional vectors computed. Thus computational cost reduction is achieved through minimising the number of *directional vectors*. Fig. 2 illustrates such computational cost reduction process through the example of two listening avatars L_1 and L_2 sharing the rendering results of the *directional vector* v_1^{220}. Instead of localising the voice of S_1 exactly according to the respective positions of L_1 and L_2, the *directional vector* v_1^{220} is distance weighted to the positions of L_1 and L_2. Perceptually, L_1 hears the voice of S_1 as emanating from the phantom position S_1' which deviates from the exact position of S_1 by β_s. Similarly, L_2 hears S_1 as coming from S_1'' with the angular deviation of α_s. Such computational sharing is permitted because the two listening avatars both have an angular deviation to v_1^{220} which is within the corresponding pre-defined thresholds called *acceptable angular errors*. We use ε_{ij} to denote the *acceptable angular error* of a listening avatar L_j with respect to a speaking avatar S_i. This concept of *acceptable angular error* was first proposed in [7]. Our computation reduction scheme is distance-governed in the sense that we assume that the *acceptable angular error* value for a particular listening avatar is determined by its distance to the corresponding speaking avatar. In the example shown by Fig. 2, the *acceptable angular error* of L_2 (ε_{12}) is smaller than that for L_1 (ε_{11}) because L_2 is closer to S_1 than L_1.

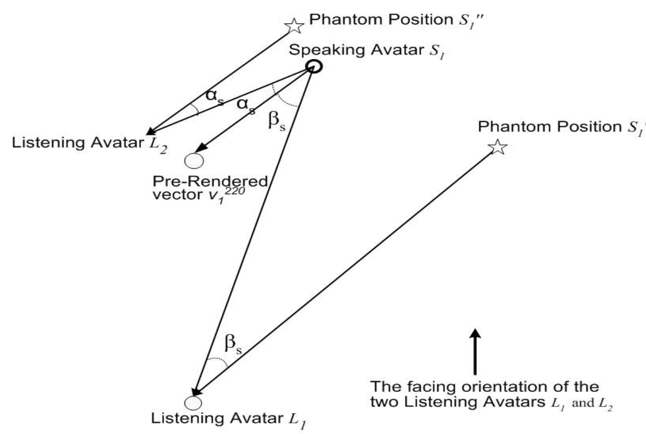

Fig. 2. Computation reduction for one instance of listening avatar

Fig. 3 shows one possible linear model of varying the *acceptable angular error* value for a given listening avatar over its distance to a given speaking avatar. This linear relationship is bounded by a maximal value denoted by ε_{max} and a minimal value ε_{min}. It must be noted that although the upper bound ε_{max} is set to a large value, but it only occurs for the few listening avatars which are on the edge of the audible range of a given speaking avatar.

It is not the intended purpose of this work to ascertain a definite *acceptable angular error* value for the virtual world or the exact model of how *acceptable angular error* would

vary. We have run our simulations over a range of *acceptable angular error* values. It is worth noting that the *acceptable angular error* values we used here are more conservative (thus offering higher rendering accuracy) than in [7] which uses the values of ε_{min} =15^o and ε_{max} =45^o. Preliminary game trials described in [7] show reasonable perceptual acceptance of the distance-governed *acceptable angular error* concept. Future trials are planned to investigate users' reactions to voice quality over varying system parameters including different *acceptable angular error* values.

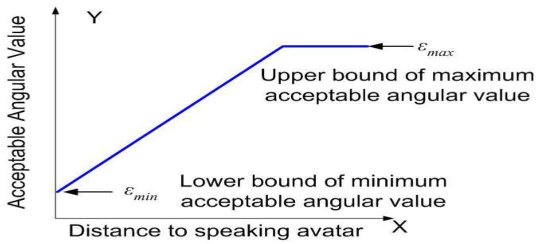

Fig. 3. The linear model of varying Acceptable Angular Error with distance

C. Model of the Virtual World used for MICE

In our simulations, the virtual world is modelled as a square area of certain size, in which avatars are distributed according to a uniform distribution. By varying the density of avatar distribution, different types of *virtual world grouping behaviour* can be simulated. A dense avatar population simulates *crowds* of users at places such as sporting stadiums and market places. A sparse avatar population simulates *loners* which are users separated by large distance and have little interactions between themselves. In this work, we analyse our distance-governed computation reduction scheme for a MICE supporting 900 avatars. Our analysis herein is conducted at a particular time instant. However, an important issue to be addressed in our future work for MICE, is to deliver MICE in a manner that can cope with avatar dynamics, for instance, the avatar movements in the virtual world and the joining/leaving of the virtual world by avatars. The size of the virtual world area was varied to change the density of our avatar population. We also assume all the avatars can hear but only half of them are speaking at any given time with a fixed audible range of 30 m.

III. The Vector Number Minimisation Model

The *vector number minimisation model* (LP1) finds for each speaking avatar, the minimal number of *directional vectors* needed to ensure that all the surrounding listening avatars can hear the speaking avatar voice within their respective *acceptable angular error* constraints. The Head Related Transfer Function (HRTF) is measured and stored as the Head Related Impulse Response (HRIR) which is the time domain representation of HRTF. HRIRs are often referred to as the HRTF filters. For each defined *directional vector*, two corresponding HRTF filters (for left and right ears respectively) are retrieved from a pre-built database [5] and then convolved (binaural synthesis) separately with the original sound source to create the left and right channel signals of that *directional vector*. Consequently, the length of the HRTF filter used determines the complexity of computing each *directional vector*. In LP1, we have used only one compact type of HRTF filter with a fixed cost of *C* equal to 128 [5].

Known variables:

Let θ_{ij}^k denote the precomputed angular deviation between a listening avatar L_j and a possible *directional vector* v_i^k , with respect to a particular speaking avatar S_i.

a_{ij} is 1 if L_j can hear S_i, 0 otherwise.

t_{ij}^k is 1if L_j can hear S_i and $\theta_{ij}^k \le \varepsilon_{ij}$, 0 otherwise

Decision variable:

x_i^k is 1 if the *directional vector* along angle *k* is created for *Si*, 0 otherwise. (1)

Objective Function Minimise:

$$\mathrm{LP1} = \sum_{i=1}^{B}\sum_{k=1}^{N} C x_i^k \tag{2}$$

Subject to:

$$\sum_{k=1}^{N} t_{ij}^k x_i^k \ge a_{ij} \quad \forall\, i, j \tag{3}$$

$$x_i^k \in \{0,1\} \;\; \forall\, i, k \tag{4}$$

In LP1, constraint (3) ensures that for each listening avatar in the audible range of each speaking avatar, there exits at least one *directional vector* that is within the appropriate *acceptable angular error* constraint. LP1 is a pure integer linear programming problem [9] but it is decomposable on a per speaking avatar basis. Such decomposability stems from the fact that the speaking avatars are all independent of each other in terms of how computation reductions are carried out for their respective groups of surrounding listening avatars. The size of LP1 is rather small on a per speaking avatar basis. The pre-processing operation can find a subset of *directional vectors* that will be used for a particular speaking avatar. This means the number of decision variables is equal to or less than *N* (the total number of possible *directional vectors* around each speaking avatar). The number of constraints equals to the number of listening avatars that can hear a particular speaking avatar. On average, this number is given by the avatar density of the virtual world.

A. The Effect of Varying Acceptable Angular Error Values on Computation Reduction

In our analytical results, the computational reductions achieved by our models are measured as a percentage reduction (*computational percentage reduction*) against the benchmark cost set by the "brute force" approach of providing MICE, where the voice of each speaking avatar is rendered according to the exact positions of all the surrounding listening avatars. At a given distance, the value of *acceptable angular error* depends on the maximal value ε_{max}. Hence for the horizontal scale of Fig. 4, we can use ε_{max} values to represent variations in the *acceptable angular error* constraint.

Fig.4 shows that the *computational percentage reduction* increases with rising ε_{max} values, for both avatar densities of 10 and the higher density of 25. This is because, with respect to a particular speaking avatar, a larger *acceptable angular error* value allows the surrounding listening avatars to share the

rendering results of fewer *directional vectors* than if a smaller *acceptable angular error* is applied. However, the rate of increase in the *computational percentage reduction* gradually slows down with the rise of *acceptable angular error* for both avatar densities studied. This effect is more evident when both avatar densities and *acceptable angular error* are higher. For instance, at density of 25 and for the ε_{max} values from 30^o to 45^o, the *computational percentage reduction* increases at 2% or lower per 5^o increase in ε_{max}. This observation is due to the fact that at higher avatar densities, there are more listening avatars around each speaking avatar and the angular separations between adjacent listening avatars are smaller. This leads to more cases of computational sharing among adjacent listening avatars. Consequently, large *computational percentage reductions* can already be achieved with a small *acceptable angular error* constraint imposed, leaving little room for improvement through relaxing the *acceptable angular error* constraint.

B. *The Effect of Varying Density on Computation Reduction*

Fig. 5 shows the computational percentage reduction over increasing avatar density for a set of five ε_{max} values between 5^o and 45^o. The trend observed is that, at a given ε_{max}, *computational percentage reduction* increases with rising avatar densities. This is due to the reason given in 3.1 of Section 3 that as density rises, there are more cases of computational sharing among the listening avatars around each speaking avatar. Moreover, the trends observed on Fig. 4 can be seen again in Fig. 5. The gaps between the curves on Fig. 5 indicate the improvements in computational percentage reduction brought by relaxing the *acceptable angular error* constraint. On the other hand, the inter-curve gap reduces steadily from one pair of curves to the next, agreeing with the previous finding that the rate of increase in computational percentage reduction gradually slows down with the rise of *acceptable angular error*. In Fig. 5, this effect is more evident at higher avatar densities as previously found on Fig. 4.

IV. The Multi-filter Cost Minimisation Model

In LP1, due to the use of one type of fixed-cost HRTF filter, computational cost reduction is achieved solely through the minimisation of the number of *directional vectors* computed. However, according to the state of art research results in the field of *HRTF localisation* [8], there is a range of HRTF filters we can employ with a varying range of rendering accuracies and computational complexities. Hence, in the multi-filter cost minimisation model (LP2), we divide the audible zone around each speaking avatar into two sub-zones. The first sub-zone is called the interactive zone which encompasses the listening avatars which are very close to the speaking avatar. The interactive zone listening avatars receive *directional vectors* rendered using high quality HRTF filters (assume to be *C* the same as LP1 for simplicity). The second sub-zone, i.e. the background zone encompasses the listening avatars situated outside the interactive zone. The background zone listening avatars receive *directional vectors* rendered using low quality HRTF filters (cost represented by C_1 set to be 12, equals to one eleventh of C [8]). In our analytic work, the radius of the interactive zones in LP2 was initially set to 5m. This radius was later varied in our analytical work to study the effect of such changes. LP2 seeks to improve further on the computational saving achieved by LP1 through finding the optimised mix of the low quality and high quality *directional vectors*.

Known variables:

e_{ij} is 1 if L_j is in the interactive zone of S_i , 0 otherwise.

f_{ij} is 1 if L_j in the background zone of S_i, 0 otherwise.

Decision Variables:

$x_{1_i}^{\ k}$ is 1 if the directional vector at angle *k* is created using the high quality HRTF filter, 0 otherwise. (5)

$x_{2_i}^{\ k}$ is 1 if the directional vector at angle *k* is created using the low quality HRTF filter, 0 otherwise. (6)

y_i^k is 1 if the directional vector at angle *k* is created using either the high quality or the low quality HRTF filters , 0 otherwise. (7)

Objective Function Minimise:

$$\sum_{i=1}^{B}\sum_{k=1}^{N} Cx_{1_i}^{\ k} + \sum_{i=1}^{B}\sum_{k=1}^{N} C_1 x_{2_i}^{\ k} \tag{8}$$

Subject To:

$$\sum_{k=1}^{N} x_{1_i}^{\ k} \geq e_{ij} \quad \forall i,j \tag{9}$$

$$\sum_{k=1}^{N} y_i^k \geq f_{ij} \quad \forall i,j \tag{10}$$

$$y_i^k = x_{1_i}^{\ k} + x_{2_i}^{\ k} \ \forall i,k \tag{11}$$

$$x_{1_i}^{\ k}, x_{2_i}^{\ k}, y_i^k \in \{0,1\} \ \forall i,k \tag{12}$$

In LP2, constraint (9) ensure that for each listening avatar in the interactive zone of each speaking avatar, there exists at lease one directional vector created with high accuracy filter that is within the corresponding *acceptable angular error* constraint. Constraint (10) ensures the all the listening avatars in the background zones of speaking avatars have at least one directional vector (high or low accuracy) within *acceptable angular error*. Constraint (11) states that the same directional vector can not be duplicated in the interactive zone and the background zone of a given speaking avatar. Note this formulation can also be extended to allow the use of more than two types of filters. Although similar to LP1, LP2 is still decomposable with respect to each speaking avatar and has reasonable computational scalability; LP2 is computationally more complex than LP1 with more binary decision variables and constraints.

Due to the space limitation and for clarity of illustrations, we have only shown in Fig. 6, the ratio difference in computational percentage reduction between LP1 and LP2 when ε_{max} is 35 degrees. All the three curves are above the ratio of 1 indicating that LP2 consistently outperforms LP1 in computational cost reduction for the three interactive zone radiuses of 5 to 15 m. The extent of this improvement by LP2 over LP1 reduces with increase in the interactive zone radius used. Decreasing the interactive zone radius allows more avatars to be rendered with the low quality but computationally efficient HRTF filter and results in lower computational cost overall than that achieved by LP1.

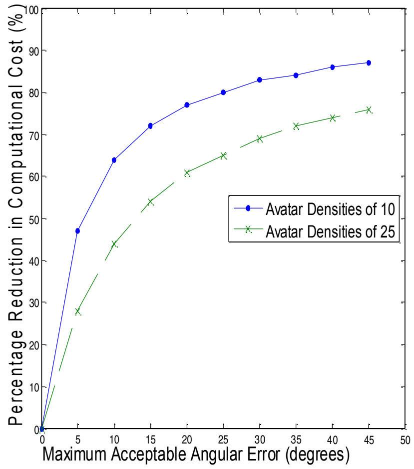

Fig. 4. The effect of varying $\varepsilon_{acceptable}$ on computational cost reduction

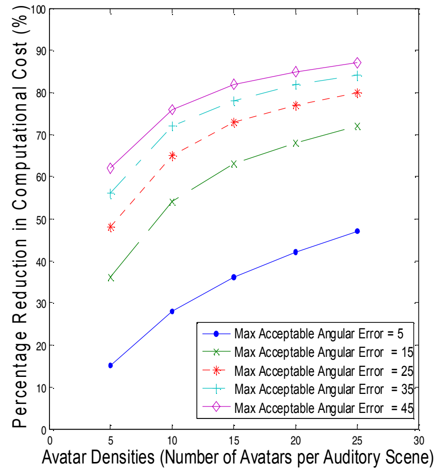

Fig. 5 The effect of avatar density on computation reduction

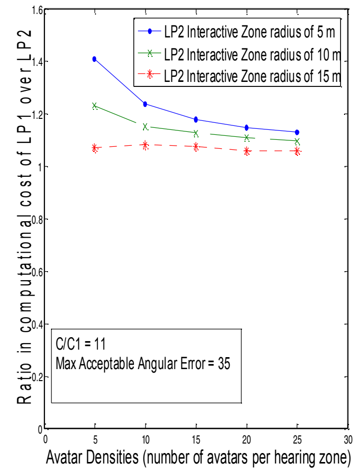

Fig. 6 The effect of avatar densities and interactive zone radius on LP2 performance

Furthermore, for both the interactive zone radii of 5 and 10 m, the ratio difference between LP1 and LP2 decreases with rising avatar densities. This is due to the fact that as found on Fig. 5 previously, at high avatar densities, large computational percentage reductions can already be achieved by LP1, leaving little room for improvement through the use of LP2.

V. RELATED WORK

None of the prior art systems reviewed thus far are scalable enough to support true Immersive Voice Communication (IVC) for a mobile Distributed Virtual Environment (DVE), especially when the number of avatars is large and or the distribution of avatars is highly dense. There are some high fidelity immersive audio rendering systems such as that described in [10] which can render multiple sound sources but not on a distributed basis for multipoint-to-multipoint IVC. Such systems often attempt to simulate sophisticated room acoustical effects which are too computationally expensive to be applied for a networked virtual environment [10]. On the other hand, the current networked voice communication systems are either text-based [11] or simple mono Voice over IP (VoIP) applications [12], neither of which can really provide truly immersive communications.

VI. CONCLUSION

In this paper, we propose two optimisation models to minimise the computational cost of providing a Mobile Immersive Communication Environment (MICE) while satisfying the rendering accuracy constraint imposed by the appropriate *acceptable angular error* values. LP1 minimises the computational cost of MICE through finding the minimal number of rendering operations (a minimal set of *directional vectors*) required to ensure that each instance of Listening Avatar in a particular MICE, is rendered within the corresponding *acceptable angular error* constraint. LP2 explores the possibility of using different types of HRTF filters to further reduce the computational cost of MICE from that achieved by LP1. Our analytical results examined the effect of various factors on the computational reductions achieved by our optimisation models. These factors include, the *acceptable angular error* values chosen, the type of HRTF filters used, and avatar densities among others. The two optimisation models proposed are decomposable and offer reasonable real-time performance. Hence they can be applied in real-time to significantly reduce the computational cost of providing a MICE system. Such low computational cost coupled with its bandwidth efficient server-based delivery architecture, provides a scalable IVC service for mobile DVE.

ACKNOWLEDGEMENT

This work is supported by the Co-operative Research Centre for Smart Internet CRC (SITCRC) and the University of Wollongong (UOW), Australia.

REFERENCES

[1] Paul Boustead and Farzad Safaei, "Comparison of Delivery Architectures for Immersive Audio in Crowded Networked Games", in *Proc. of the 14th ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV),* Ireland, 16-18 June, 2004.

[2] J. Bolot and F. Parisis, "Adding Voice to Distributed Games on the Internet", in *Proc. of the Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies,* pages 480-487, 1998.

[3] M. Radenkovic, C. Greenhalgh, S. Benford, "Deployment issues for multi-user audio support in CVEs," in *ACM Symposium on Virtual Reality Software and Technology*, 2002, pp. 179-185.

[4] MMOGCHART.com, http://www.mmogchart.com/ (9 Aug. 2005).

[5] Durand R. Begault, *3-D Sound for Virtual Reality and Multimedia*, Academic Press Professional, Cambridge, MA, USA, 1994.

[6] Thomas Funkhouser, Patrick Min and Ingrid Carlbom, "Real-Time Acoustic Modelling for Distributed Virtual Environments", in *Proc. of the 26th annual conference on Computer graphics and interactive techniques,* pp. 365-374, 1999.

[7] Paul Boustead, Farzad Safaei and Mehran Dowlatshahi, "DICE: Internet Delivery of Immersive Voice Communication for Crowded", in *Proc. of IEEE Virtual Reality (VR) 2005,* Bonn, Germany, Mar. 12-16, 2005.

[8]. Jyri Huopaniemi and Matti Karjalainen, "Review of Digital Filter Design and Implementation Methods for 3-D sound", in *Proc. of AES 102nd convention*, Preprint 4461, Mar. 1997.

[9] Ronald L. Rardin, *Optimization in Operations Research*, Prentice Hall, Upper Sandle River, N.J. USA, 1998, pp. 561-562.

[10] Martin Naef, Oliver Staadt, Marjus Gross, "Spatialised Audio Rendering for Immersive Virtual Environments", in *Proc. of the ACM symposium on Virtual reality software and technology,* pp 65-72, 2002.

[11] ICQ Inc, ICQ Home Page, http://www/icq.com, (9 Aug. 2005).

[12] TeamSpeak, The Team Pay Engine, http://www.teamspeak.org, (9 Aug. 2005).