# BRIDGING THE GAP BETWEEN FIXED AND MOBILE ACCESS TO A LARGE-SCALE NVE INCORPORATING BOTH AUDIO AND VIDEO

Peter Quax Maarten Wijnants Tom Jehaes Wim Lamotte
Expertise Center for Digital Media
Limburgs Universitair Centrum
Universitaire Campus
B-3590 Diepenbeek, Belgium
{peter.quax, maarten.wijnants, tom.jehaes, wim.lamotte}@luc.ac.be

**ABSTRACT**

Networked Virtual Environment technology has pushed the boundaries of a number of currently popular applications on the Internet. On-line virtual communities and large-scale on-line games are representative examples of the commercial appeal of the research in this area. However, these applications have not yet made it into the mobile 'scene'. In general, when analyzing the current market for mobile applications, the type of applications that can be distinguished is twofold. On the one end, existing applications, already in use on wired networks, are being extended to include mobile devices as deployment platform. On the other hand, entirely novel applications, targeted specifically towards some of the peculiar aspects of mobile use, have seen the light of day. Deploying networked virtual environments (NVEs) on mobile devices should be considered an opportunity to mix these two application types. In this paper, we will show how NVE technology can be extended towards mobile use while retaining backward compatibility with fixed network setups We will also offer some insight on novel applications of NVEs on this newly targeted platform.

**KEY WORDS**
Networked Virtual Environments, Networking, Mobile Computing, Technology Adaptation.

## 1 Introduction and related work

In recent years, we have been witnessing an explosive growth in the number of commercially deployed Networked Virtual Environment (NVE) applications. This is especially true in the field of computer entertainment considering the popularity of Massive Multiplayer On-line Role-Playing Games (MMORPGs) [3] and the Virtual Interactive Communities [11][23]. Subscription numbers vary from a few hundred to several thousands, turning it into a major source of revenue for publishers and dot-com businesses alike.

The same growth can be observed in the field of mobile devices, especially mobile phones. Simple applications on these devices have shown to be able to yield major profits for providers, e.g. Short Message Services (SMS), ringtone downloading and tele-voting on quiz shows. Most of these applications are in fact fairly simple to implement on an existing network architecture.

Some effort has been put into providing networked virtual environment technology to mobile device users. In this context, we should mention the graphically simple games that are currently available on mobile phones and personal digital assistants (e.g. Mission 3D [8] and Badlands [24]). These examples all rely on ad-hoc BlueTooth short-range communication networks as underlying network technology. While these examples do allow for a number of players to participate in a single virtual environment, they can not be regarded as serious fully-fledged NVE applications. This is primarily due to the fact that they only support a limited number of simultaneously connected users and, secondly, are targeted towards a very specific type of mobile device.

A more serious effort was made this year by Sega with Pocket Kingdom [22], which is referred to as the worlds' first MMORPG suitable for deployment on a specific type of mobile phone (the Nokia N-Gage gaming deck). The N-Gage Arena service [15], as Nokia calls it, provides for users to participate in world-wide on-line gaming through a GPRS connection. This is combined with a web-based service that provides tracking of the activities/skills/experience of each single user. The popularity of this type of application on mobile devices remains, as of yet, to be determined. However, a seriously limiting factor for potential customers is the fact that these games are limited to mobile-phone use only (and, in fact, only those of the Nokia brand). Should applications, such as World of Warcraft [3], which has proved to be immensely popular in previous months, be playable on a mobile platform, it is likely that they will provide an even deeper immersion in the on-line virtual environment of the MMORPG.

In this paper, we will discuss some of the networking issues and problems to be resolved in connecting both wired and wireless devices to the same large-scale networked virtual environment application. Therefore, we will consider the extension of our in-house developed NVE architecture. For a general introduction to the architecture, we would like to refer to our previous work in [9], [10] and

[17].

This rest of this paper is organized as follows : Section 2 provides a short introduction to the wired-only version of the NVE architecture. In Section 3, we will provide, in general terms, some of the underlaying networking technology needed for mobile access to NVE applications. Section 4.1 discusses the easiest and most obvious way of extending the architecture to allow for wireless mobile device access. In Section 4.2, we will discuss a solution for long range access to an existing NVE. Section 4.3 extends on this solution by including mobile mini-lans, providing a more general architecture that sets off some of the limitations of the previous approaches. Section 5 provides some conclusions and pointers for future work.

## 2 Wired NVE access

To clarify some of the decisions made in the process of extending the framework towards mobile access, as described in this paper, we will provide a short overview of the architecture of the ALVIC framework in this section. For more detailed information, we would like to refer to the more detailed descriptions in [9], [10] and [17]. It should be noted that the primary goals taken under consideration when designing the framework were scalability, adaptability and extensibility. As all of these features have their repercussions when considering mobile access, we will describe them in some more detail in the following paragraphs.

According to the philosophy of the system, the entire virtual world is divided into a number of regions, each of which is associated with its own unique IP multicast address. The system only requires a client to send data to the multicast group of the region it is located in at each given time. To determine which data to receive, we have employed the concept of an Area-Of-Interest (AOI). At each given time, the area of interest contains a list of multicast groups that are of 'interest' to the users, e.g. are in his/her line of sight. Depending on the regions that are in a clients AOI, multicast groups are joined and left as the client moves around the world. This way, we achieve a highly scalable system in which clients can easily adapt the incoming data flow depending on available bandwidth or processing power. Reduction of downstream bandwidth is achieved by scaling down the number of subscribed multicast addresses. Distribution of action data in this scheme is automated and does not require any server intervention. It is important to note that the upstream bandwidth use of any given client is never influenced by the active number of clients in its AOI. The architecture as presented here was extensively tested for its scalability properties, and yielded satisfactory results. Results of the different testing strategies and scenarios can be found in [19].

The architecture includes a server setup that can be considered minimal compared to a client-server based system. On the one hand, the *master servers* take care of accounting, administration and forwarding. On the other hand, the *game servers* are responsible for tracking the pool of multicast addresses associated with world regions and informing the clients of those addresses when requested. This setup is easily distributable because of the light load posed on each of the servers, combined with the fact that the 'world' is, by nature of the architecture, already divided in a number of regions.

The extensibility of the framework comes into play when data streams other than pure positional/state information were added to the architecture. In this case, we added the ability for inter-person communication using 'video-based avatars'. These representations of the users in the virtual world consist of a generic geometrical form, on which a grabbed camera image of the individual user is displayed in real-time.

As soon as large-volume, high-bandwidth requiring data streams such as audio/video were introduced into the system, the need for adaptability became clear, as not all clients were able to transmit and/or receive data from all clients in the virtual world. We adapted the original architecture to include 'video area of interest' managers that were able to decide autonomously what streams to accept and reject. We chose to let each video client transmit his/her video streams in three different qualities, each to a separate IP multicast group, associated with a region in the virtual world (as in the original architecture for positional/state information). By subscribing to a specific set of multicast regions, according to the available downstream bandwidth, optimal use of the downstream link can be achieved. At the same time, the upstream link does not have to cope with bursty traffic, as the number of viewers does not affect the upstream bandwidth requirements. Test results of these additions are described in [16] and [18].

## 3 Networking Technology for Mobile Access

Almost all current PC based network applications are based on the IPv4 protocol stack. While this has, and will for some years, suffice for most applications, a gradual switch to IPv6 is anticipated in the near future. The need for this migration is clear when considering the fact that even the smallest of devices will be connected to some network or another (most likely the Internet) [2]. Besides the greatly enlarged address space, IPv6 also features extensions over IPv4 that make it suitable for deployment on mobile platforms, such as provisions for roaming and handovers [6][25]. Mobile extensions for IPv4 have been devised [1], but will probably be made superfluous by the advances made in the IPv6 protocol stack. Most of the research into mobile applications is based on these next generation protocol stacks [12][20]. Wired networked applications are based on the fact that computers are stationary devices and can therefore rely on the assumption that they have a fixed and permanent connection to the network/Internet through a single uplink. These facts are no longer true when considering the context of mobile devices. The traditional fixed connection is replaced by a continuously changing network layout, introducing all sorts of

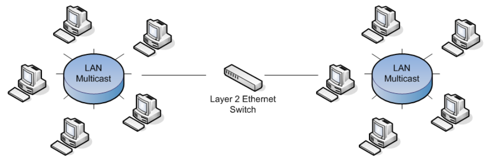


Figure 1. Existing Architecture for wired NVE access

anomalies such as packet loss, unpredictable delay/jitter values, temporary loss of connection, There may even no longer be a direct uplink to the Internet as such; moreover, a device may have to dynamically interconnect with other devices in its neighborhood in order to find a route to its destination. This creates so-called ad-hoc networks, which has been a major research topic in the last few years [14]. Special interest has been given to packet routing in these networks [21].

While multicasting was available for IPv4, its implementation and market penetration was not a great success. Many of the early network devices did not even have built-in support for multicasting. Even today, there are many practical issues to be overcome when deploying a multicast-enabled application on the Internet. Routers may block multicast traffic on specific addresses, firewalls will not allow users to join multicast groups by blocking IGMP (Internet Group Management Protocol) traffic, Ethernet switches may have support for IGMP snooping but for older versions of the protocol only, etc... Multicast extensions for mobile IPv4 are even more cumbersome [5]. Instead of using a separate address range, IPv6 incorporates multicast as an inherent part of the addressing scheme, which allows for a much greater range of possible multicast groups to be supported. Modern hardware that supports IPv6 also has provisions for multicast built-in, a turnover that is triggered by broadcasting companies that wish to distribute e.g. television programs over IP based networks.

## 4 Architectural Extensions Towards Mobile Access

New technologies such as 802.11 WLAN [7] and BlueTooth [4] are replacing the traditional Ethernet protocol for wired networks. These technologies are being integrated in a wide range of mobile devices such as PDAs and SmartPhones. In the following subsections, we will look into possible extensions of our existing wired NVE framework to enable access for mobile clients using these kinds of devices.

### 4.1 Extensions for short-range mobile access

The most obvious way of extending the existing architecture (see fig 1) towards mobile access is depicted in fig 2. A number of wireless access points, each connected

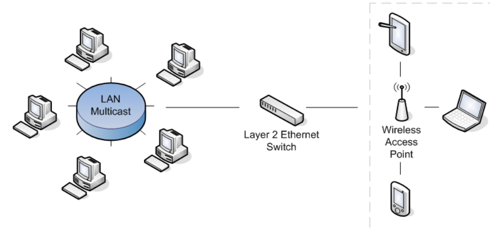


Figure 2. Extended architecture for short range mobile access

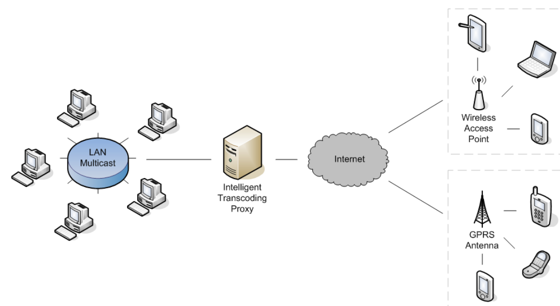


Figure 3. Proposed architecture for long range mobile access

to the LAN, forward traffic over a wireless connection to the mobile devices. While this seems an obvious approach to tackle the connectivity problems, there are some issues to be resolved. First of all, bandwidth, and more specifically throughput, over a wireless (radio) link is limited when compared to a Fast Ethernet LAN. This leads to problems when dealing with an NVE in which a large number of users are employing video-based avatars. These issues can, up to a certain degree, be dealt with in a relatively straightforward fashion. As the framework makes use of multicast for distributing both event and audio/video data, we can already filter out a great deal of unnecessary traffic at protocol layer 2 (link layer). The majority of manageable ethernet switches nowadays are able to 'snoop' multicast traffic that is sent over the wire. This is accomplished by eavesdropping on the IGMP traffic that is communicated between clients. By intercepting both multicast 'join' and 'leave' messages, a switch is able to determine what groups should be forwarded over each of its links. If such a multicast-enabled switch were to be used to connect each of the access points to the wired LAN, a large number of users may be connected directly to the multicast-enabled LAN. The major downside of this approach is the limited range of these WLANs.

## 4.2 Enhanced architecture for long-range mobile access

To enable mobile access to NVE applications over truly large distances, a number of connection methods should be considered. At short range, WLAN 802.11 would be an obvious candidate if such an infrastructure is at hand. In other cases, a GPRS or similar connection may be the only way for a mobile device to connect to the virtual environment. It should however be clear that in these cases, it is exceedingly difficult to provide users with a view of the virtual environment comparable to the one that is available when connected directly to the multicast-enabled LAN.

This is why we propose to insert an intelligent proxy at the edge of the LAN which provides a number of services specifically targeted towards mobile users. The resulting architecture is depicted in figure 4.2. In case of the sample NVE architecture, we envision the proxy fulfilling two major tasks. First of all, the proxy should act as a multicast-to-unicast gateway for remote mobile clients. As there is very little support for multicasting in the Internet at present, traffic being sent on the multicast-enabled LAN should be forwarded towards each remote mobile client through unicast connections. However, simply duplicating all multicast network traffic that resides on the LAN does not scale well with a growing number of connected users, neither is it suited for low-bandwidth connections such as GPRS. For these reasons, we believe the proxy should have an extensive knowledge of both the application it is serving as well as the network infrastructure. Based on its compound application and network awareness, the proxy can then intelligently decide which multicast traffic should be unicasted to each individual remote client.

Secondly, we envision that the proxy will also be acting as a transcoder of audio and video streams sent out by wired users. Clients connected directly to the LAN will most likely be exchanging high-quality audio and video streams whose bandwidth requirements by far exceed the available throughput of a long-range wireless link. The proxy should therefore be able to transcode multimedia streams in real-time to lower-quality versions that better fit the bandwidth limitations imposed by wireless connections. Again, the proxy's application and network awareness will play an important role in deciding which multimedia streams should be transcoded and to what quality. However, the capabilities of the mobile device the multimedia stream is destined for should also not be neglected when making transcoding decisions. For example, it is pointless to send a high resolution encoded video to a mobile phone user due to the limited screen size of these devices. As a result, the proxy should also gain device awareness by gathering information regarding the capabilities of client devices.

To summarize, it is the responsibility of the proxy to offer mobile users a maximal Quality of Experience (QoE) without ever exceeding their available downstream bandwidth. To achieve this goal, the proxy will have to intelligently limit the data flow towards mobile clients and transcode multimedia streams to lower qualities as needed. For a more detailed description of our proxy solution, we would like to refer the interested reader to [13].

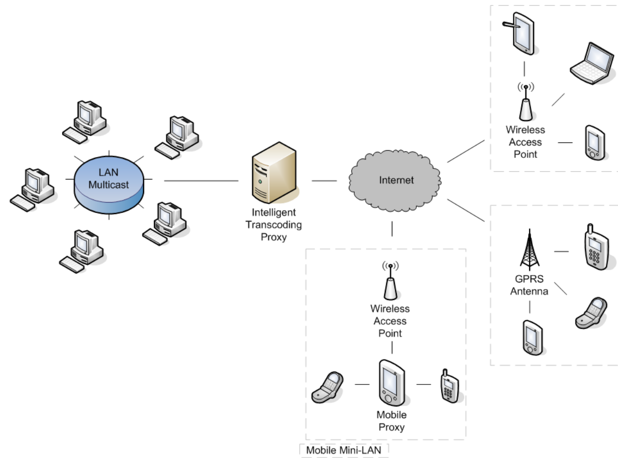


Figure 4. Proposed architecture including Bluetooth-based short-range mini-LAN

## 4.3 Ad-hoc enhancements for mobile mini-networks

As detailed above, placing an intelligent proxy at the edge of the multicast-enabled LAN, combined with WLAN and GPRS mobile connectivity, enables access to the NVE application for a large number of devices. However, there are still some that may not be covered due to, for example, lack of availability of long-range radio functionality in devices. Our third proposed extension to the architecture tries to solve some of these issues by using ad-hoc network technology between mobile devices in close neighborhood, forming local meshes.

In this scenario, a powerful mobile device will perform part of the proxy functionality that was described in section 4.2. By maintaining BlueTooth connections between this device and each of the clients in its vicinity, we will be able to transmit the necessary data over these short-range wireless channels, forming a local mesh or virtual mini-LAN (see fig 4.3). The 'mobile proxy' is responsible for providing the uplink to the rest of the network, either through a GPRS or WLAN connection. In the case of a GPRS connection, the available bandwidth is severely limited, so the mobile proxy should instruct the fixed proxy to discard all video streams at the edge of the multicast-enabled LAN, leaving only the positional and state information available for the mobile clients in range of the mobile proxy.

Given the roadmap for BlueTooth development presented in [ref], we may even consider passing video streams along the ad-hoc BlueTooth network when a high-bandwidth WLAN uplink is available. This can be achieved because multicasting data along BlueTooth connected hosts will be supported for up to 7 devices. Some caution is needed however, as the achievable data rates are currently unknown due to lack of available hardware.

## 5 Conclusions and future work

We have presented three different strategies to extend access to networked virtual environments towards mobile devices. The first scenario makes use of existing hardware and software to achieve limited range accessibility, by simply adding wireless access points to a switched LAN setup. While easiest to implement, this scenario has a range of drawbacks, which is the primary reason why the second scenario, the inclusion of intelligent proxies in the network, may be seen as a more complete solution to the problem. It not only enables longer range communications but also adapts streams for each specific client to accommodate those with lower-bandwidth links (such as GPRS). Finally, the third strategy looks into the near future with the proposition of mobile mini-lans that form in an ad-hoc fashion. This enables both transmission of data between users (as in a wired LAN setup) and enables connectivity to wired networks using mobile proxies.

The first two proposed strategies are currently implemented into the ALVIC architecture and scalability is being tested. The load on the intelligent proxy while transcoding streams for a high number of users is of particular interest, which is why measurement test are being undertaken.

The third strategy, incorporating mobile proxies and mini-lans will be added in a second stage. Because multicasting on bluetooth communication channels is unavailable in current hardware (but foreseen in the bluetooth roadmap), we will resort to simulations for this type of setup.

Besides the positional, state and audio/video traffic, there are some other streams present in a networked virtual environment, which were not discussed in this paper. These consist primarily of geometrical and/or animation-related information. Work is currently undertaken to analyze this traffic and adapt it for use on mobile devices.

## 6 Acknowledgements

Part of this research was funded by IBBT ( Institute for BroadBand Technology ) of Flanders and the Flemish government, We would like to thank all members of the NVE research group at EDM for their support in developing the technology as presented in this paper.

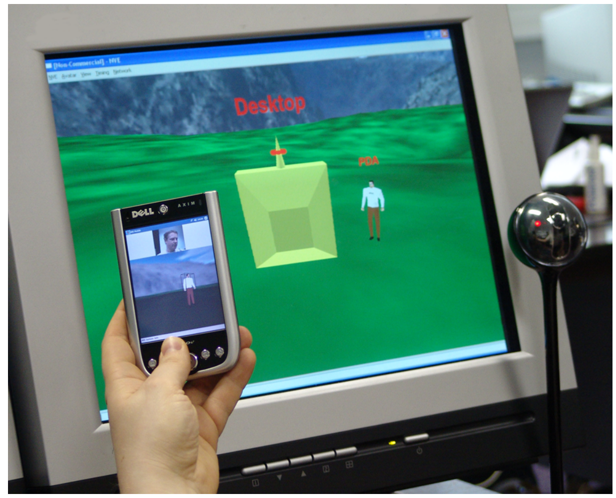


Figure 5. Mobile access to our NVE using a PDA. A user's webcam video is streamed to the mobile client on request.

## References

[1] Ip multicast extensions for mobile internetworking. In *Proceedings of the Fifteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, 1996.

[2] O. Ardaiz, F. Freitag, and L. Navarro. On service deployment in ubiquitous computing. In *Proceedings of the 2nd International Workshop on Ubiquitous Computing and Communications*, 2001.

[3] Blizzard. *World of Warcraft*. World Wide Web, `http://www.worldofwarcraft.com`.

[4] Bluetooth Consortium. *The official bluetooth website*. World Wide Web, `http://www.bluetooth.com`.

[5] Vineet Chikarmane, Carey L. Williamson, Richard B. Bunt, and Wayne Mackrell. Multicast support for mobile hosts using mobile ip : Design issues and proposed architecture. *Mobile Networks and Applications*, 3(4):365–379, 1999.

[6] Eric CARMS. *IPv6: answers to your questions*. White paper on Ipv6 Forum, `http://www.ipv6forum.org/navbar/documents/6WIND-Ipv6-answers-v2.0.pdf`.

[7] IEEE. *802.11 standard*. World Wide Web, `http://standards.ieee.org/getieee802/`.

[8] In-Fusio. *Mission 3D*. World Wide Web, `http://www.in-fusio.com`.

[9] J. Liesenborgs, P. Quax, et al. Designing a virtual environment for large audiences. *Lecture Notes in Computer Science*, Information Networking, Wireless Communication Technologies and Network Applications:585–595, 2002.

[10] J. Liesenborgs, P. Quax, et al. Designing a virtual environment for large audiences. In *Proceedings of the 16th International Conference on Information Networking (ICOIN)*, pages 3A–2.1–3A–2.10, 2002.

[11] Linden Labs. *SecondLife*. World Wide Web, http://www.secondlife.com.

[12] Keith Mitchell, Duncan McCaffery, George Metaxas, Joe Finney, Stefan Schmid, and Andrew Scott. Six in the city: introducing real tournament a mobile ipv6 based context-aware multiplayer game. In *Proceedings of the 2nd Workshop on Network and System Support for Games*, 2003.

[13] Patrick Monsieurs, Maarten Wijnants, and Wim Lamotte. Client-controlled qos management in networked virtual environments. In *Proceedings of the Fourth International Conference on Networking*, pages 268–276, 2005.

[14] National Institute of Standards and Technology. *Wireless Ad-Hoc Networks Bibliography*. World Wide Web, http://w3.antd.nist.gov/wctg/manet/manet_bibliog.html.

[15] Nokia. *N-Gage Arena*. World Wide Web, http://arena.n-gage.com.

[16] Peter Quax, Chris Flerackers, Tom Jehaes, and Wim Lamotte. Scalable transmission of avatar video streams in virtual environments. In *Proceedings of the 2004 IEEE International Conference on Multimedia and Expo (ICME 2004)*. IEEE, 2004.

[17] Peter Quax, Tom Jehaes, Pieter Jorissen, and Wim Lamotte. A multi-user framework supporting video-based avatars. In *Proceedings of the 2nd workshop on Network and system support for games*, pages 137–147. ACM Press, 2003.

[18] Peter Quax, Patrick Monsieurs, Tom Jehaes, and Wim Lamotte. Performance evaluation of client-side video stream quality selection using autonomous avatars. In *Proceedings of the 2004 International Conference on Advances in Computer Entertainment Technology (ACE 2004)*, pages 251–256. ACM Press, 2004.

[19] Peter Quax, Patrick Monsieurs, Tom Jehaes, and Wim Lamotte. Using autonomous avatars to simulate a large-scale multi-user networked virtual environment. In *Proceedings of the 2004 International Conference on Virtual-Reality Continuum and its Applications in Industry (VRCAI2004)*, pages 88–94. ACM Press, 2004.

[20] Sebastin Matas Riera, Oliver Wellnitz, and Lars Wolf. A zone-based gaming architecture for ad-hoc networks. In *Proceedings of the 2nd Workshop on Network and System Support for Games*, 2003.

[21] S. Corson and J. Macker. *Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations.* IETF Draft 1998.

[22] Sega. *Pocket Kingdom*. World Wide Web, http://www.pocket-kingdom.com.

[23] There Inc. *There*. World Wide Web, http://www.there.com.

[24] VivaLaMobile. *Badlands*. World Wide Web, http://www.badlands.com.au.

[25] Wolfgang Fritsche and Florian Heissenhuber. *Mobile Ipv6. Mobility Support for the Next Generation Internet*. White paper on Ipv6 Forum, http://www.ipv6forum.com/navbar/papers/MobileIPv6_Whitepaper.pdf.