Samen aanjagen van vernieuwing

Scaling XR with new wireless technologies

in education and research

Auteur(s): Jeroen Boots (XR Zone, TU Delft), Luuk Goossen (XR Zone, TU Delft), Tim Kok (SURF)
Datum: Januari 2024

Europese Unie
Europees Fonds voor Regionale Ontwikkeling

Executive Summary

In this report, we explore how the increased bandwidth and low latency of new wireless technologies (5G and WiFi 6) can be used to scale eXtended Reality (XR) in education and research. We also investigate the current network challenges of scaling XR.

After an overview of the components necessary for scaling XR in research and education, the current networking challenges in XR, and the potential benefits of 5G and WiFi 6, we discuss two test cases we conducted at the Do IoT Fieldlab: 1) XR remote rendering using 5G; and 2) XR multiplayer applications on a 5G network.

Based on our two test cases, we conclude that 5G networks likely provide sufficient bandwidth and low latency for XR applications in education and research. Nevertheless, we expect the actual usage of 5G in XR to be limited, at least in the short to medium term. The adoption of 5G in XR in the long term will depend on future developments of XR hardware and software, particularly in Augmented Reality (AR).

Another general finding from our tests is that the current state of XR software and tooling is not quite ready to support usage outside of a home environment. Indeed, XR software and tooling are probably the biggest bottlenecks in scaling XR for education and research.

Our tests also showed that the locality of the XR applications can be an important factor when designing an XR setup. When performance is critical, we recommend hosting XR applications for remote rendering and multiplayer on edge locations, i.e., on campus or in nearby data centres.

Regarding XR streaming, we expect streaming to take a larger part of the total market in the future. Although our report considered mobile network technologies and (public) cloud, we want to highlight that the first step when developing an XR application should be optimizing the XR application itself.

The main finding from our multiplayer tests was that communication between players is a problem in the context of different private networks. A solution to this problem is to use a central server that takes care of the communication between players.

How 5G mobile networks will shape the XR landscape in the long term remains to be seen, but the opportunities that 5G and (soon) 6G provide in combination with edge and cloud technologies are seemingly endless. In the meantime, there are plenty of challenges that need to be solved to start scaling XR in education and research.

Samenvatting

In dit rapport onderzoeken we hoe de verhoogde bandbreedte en lage latency van nieuwe draadloze technologieën (5G en WiFi 6) gebruikt kunnen worden om eXtended Reality (XR) op te schalen in onderwijs en onderzoek. We onderzoeken ook de huidige netwerkuitdagingen van het opschalen van XR.

Na een overzicht van de componenten die nodig zijn voor het opschalen van XR in onderzoek en onderwijs, de huidige netwerkuitdagingen in XR en de potentiële voordelen van 5G en WiFi 6, bespreken we twee testcases die we hebben uitgevoerd in het Do IoT Fieldlab: 1) XR-rendering op afstand met behulp van 5G; en 2) XR multiplayer-toepassingen op een 5G-netwerk.

Op basis van onze twee testcases concluderen we dat 5G-netwerken waarschijnlijk voldoende bandbreedte en lage latency bieden voor XR-toepassingen in onderwijs en onderzoek. Toch verwachten we dat het daadwerkelijke gebruik van 5G in XR beperkt zal zijn, in ieder geval op de korte tot middellange termijn. Het gebruik van 5G in XR op de lange termijn zal afhangen van toekomstige ontwikkelingen van XR-hardware en software, met name op het gebied van Augmented Reality (AR).

Een andere algemene bevinding uit onze tests is dat de huidige staat van XR-software en -tooling nog niet helemaal klaar is voor gebruik buiten de thuissituatie. XR-software en -tooling zijn waarschijnlijk de grootste knelpunten bij het opschalen van XR voor onderwijs en onderzoek.

Onze tests toonden ook aan dat de lokaliteit van de XR-applicaties een belangrijke factor kan zijn bij het ontwerpen van een XR-opstelling. Als performance van cruciaal belang is, raden we aan om XR-applicaties voor rendering en multiplayer op de edge locaties te hosten, d.w.z. op de campus of in datacenters in de buurt.

Wat betreft XR-streaming verwachten we dat streaming in de toekomst een groter deel van de totale markt in beslag zal nemen. Hoewel ons rapport mobiele netwerktechnologieën en (publieke) cloud in beschouwing nam, willen we benadrukken dat de eerste stap bij het ontwikkelen van een XR-applicatie het optimaliseren van de XR-applicatie zelf zou moeten zijn.

De belangrijkste bevinding van onze multiplayertests was dat communicatie tussen apparaten een probleem is in de context van meerdere privé netwerken. Een oplossing voor dit probleem is het gebruik van een centrale server die de communicatie tussen spelers verzorgt.

Hoe 5G mobiele netwerken het XR-landschap op de lange termijn zullen vormgeven, valt nog te bezien. Tegelijkertijd zijn de mogelijkheden die 5G en (binnenkort) 6G bieden in combinatie met edge- en cloudtechnologieën eindeloos. Ondertussen zijn er nog genoeg uitdagingen die moeten worden opgelost om te beginnen met het opschalen van XR in onderwijs en onderzoek.

Table of Contents

1 Introduction

The use of eXtended Reality (XR)[1] in education and research is becoming more commonplace. Where initial XR experiments were limited to a handful of devices, soon classrooms of 50 to 200 students will need to be serviced simultaneously. This new scale of XR in education and research requires changes in device management, application deployment, and networking and infrastructure, among other things. In this report, we explore, on the one hand, how the increased bandwidth and low latency of new wireless technologies (5G and WiFi 6) can be used to scale XR, both theoretically and practically. On the other hand, we investigate the current network challenges of scaling XR, specifically in the context of remote rendering and multiplayer applications.[2] As many XR Head-Mounted Displays (HMDs) use wireless technologies, currently WiFi and in the future 4G and 5G, the possibilities and pitfalls of new wireless technologies will determine the future shape of XR.

The structure of this report is as follows. In section 2 of this report, we will discuss the components necessary for scaling XR in research and education, placing additional emphasis on the networking challenges (section 2.1). In section 3, we will discuss the opportunities new mobile technologies bring to the field of XR, followed by an overview of the hardware and software that is currently available in section 4. In section 5, we will discuss tests we have done with 5G at the Do IoT Fieldlab. First, we will discuss the tests we did to explore remote rendering using 5G (section 5.1), and then our tests of multiplayer applications on a 5G network (section 5.2). While the main topic of this section concerns the 5G tests, we will inevitably touch upon some of the other factors of scaling XR, as these are closely intertwined. Finally, in section 6, we will summarize our main findings and discuss potential next steps.

This report and the 5G tests have been the result of a collaboration between the XR Zone, part of TU Delft's NewMedia Centre, and the Do IoT Fieldlab (supported by the European Development Fund). The XR Zone has many years of experience with XR in education and research, both in facilitating the use of XR and in developing their own XR applications and infrastructure. Currently, the XR Zone is developing their vision on how to scale XR at the TU Delft in the coming years. The Do IoT Fieldlab is set up by the TU Delft, TNO, MCS, and SURF to provide a testing environment to demonstrate and explore the capabilities of 5G.

[1] "Extended reality (XR) is a catch-all term to refer to augmented reality (AR), virtual reality (VR), and mixed reality (MR). The technology is intended to combine or mirror the physical world with a "digital twin world" able to interact with it." (Wikipedia)
[2] Although, in section 2, we discuss several components of scaling XR for education and research, the focus of this report will be on networking.

2 Scaling XR for research and education

As mentioned in the introduction, the XR Zone is currently working on and evaluating different strategies for scaling up XR in academic environments. These efforts can be divided into three categories: Development, Infrastructure, and Use.

Development	**Infrastructure**	**Use**
Developer Tools	Mass Device Management	Remote Control
Templates	Data Storage	Quick Setup
Plugins	SSO	Intuitive UX & Operation
	Networking	Multiplayer
	Remote Rendering	

The development category includes various tools to speed up the development of advanced XR applications. This category consists of templates for commonly recurring types of projects, plugins, and common tools to unify and standardise the interfaces and interactions between applications and automate various parts of the development workflows.

The infrastructure category includes the storage, delivery, support, and services surrounding deployed XR applications, and hardware. For example, mass device management and automatic content delivery and updates, storing application data securely, privately and with back-ups, single-sign-on (SSO) implementations for identifying users, networking between instances of an application (multiplayer), networking between an application and its support services, and the remote rendering of XR content.

Finally, the use category encompasses everything the teachers, researchers, students, and other users of the application will interact with. This includes having an intuitive and easy-to-understand user interface for users of various backgrounds, improving overall user experience, and reducing common side effects like nausea, fatigue, and headaches. Furthermore, efforts are being made regarding remote controlling XR applications, making it easier for teachers, researchers, and other observers to work with large amounts (50+) of participants at the same time. Another area of interest is the time and ease with which XR experiences can be set up and started, reducing the effort of human resources needed to provide XR to large groups of people.

Currently, XR applications and devices are mostly used within controlled environments. Usually, these are indoor labs with all the necessary equipment already in place. Part of the scaling up strategy is making it easier to deploy XR experiences across campus. This means applications should be able to run reliably using easy-to-move setups and campus network infrastructure. To accommodate this, standalone XR HMDs such as the Meta Quest 2 or Pico Neo 4 Enterprise can be used to run XR applications without the need for heavy and expensive computers. For more computationally heavy content, graphics could be rendered remotely and streamed to the destination HMDs as we will discuss in this report.

Moreover, with recent developments in Augmented Reality (AR) and Mixed Reality (MR) devices, various demonstrations, and tradeshows, and researchers requiring on-location experiments, XR applications might need to be run on locations outside of university campuses. In these situations, wired or high-end (Wi-Fi 5 or Wi-Fi 6) wireless internet connections are not always a given. In these situations, 5G could be a potential candidate to ensure a consistent system without the need for location-specific configurations.

2.1 Networking and XR

In this section, we will first give an overview of the state of networking in XR, then we will discuss the different networking contexts and how these can be challenging for current XR applications.

2.1.1 State of networking in XR

Virtually all XR software is built for a WiFi network at home. Many applications, such as Air Link, rely on the Universal Plug and Play (uPnP)[3] protocol for device discovery. This is the same protocol that is used for connecting to printers on a home network and casting media to TVs. Additionally, XR applications might also depend on multicast[4] and uncommon network ports. By contrast, many applications do not support IP-based discovery of devices.[5] Generally, there is no technical documentation for XR devices and software, and it can be difficult to figure out what the network requirements of an XR application are because most of the software and tools are proprietary. Finally, there are currently no XR HMDs with a built-in 5G modem.

2.1.2 Network options and challenges

The available network options and challenges depend on the context in which you are.

2.1.2.1 Lab

The lab is a part of the campus in which we have more control over the networking. This could be because there are no firewalls (or other restrictions) blocking traffic within the lab's network or because we have our own networking hardware to set up a private network. (Note that if you do not have this level of control over the network in your lab, you are probably operating in the campus context below.)

In this context, we generally do not encounter many problems with XR HMDs and applications because we can closely mimic the home environment. As most of the applications are specifically designed with the home environment in mind, we do not run into many networking issues in the lab. The lab network is generally considered safe and there is no incoming external traffic because we do not host applications in the lab for external use. Finally, in the lab it is possible to connect an XR HMD with a cable to a computer which circumvents a lot of networking issues, but, on the other hand, is not very scalable or flexible.

2.1.2.2 Campus

A campus network is characterized by network segmentation (i.e., there are several networks instead of one) and policy enforcement (e.g., firewalls and network isolation). In this context, we encounter challenges with XR networking. For example, most campus networks will have

[3] https://en.wikipedia.org/wiki/Universal_Plug_and_Play
[4] https://en.wikipedia.org/wiki/Multicast
[5] Some applications, such as Virtual Desktop, solve this discovery problem by introducing a central server to which all HMDs register first and that matches up HMDs with each other by exchanging IP information.

disabled uPnP for (legitimate) security reasons and non-standard ports will be blocked. It might also be difficult to reach HMDs that are located in different networks on the campus (e.g., connectivity between an XR HMD on eduroam and a workstation connected via an ethernet cable).[6] This might be because traffic is blocked by a policy or because the application does not work with IP-based routing. An (on-premise) datacentre could be part of the campus network and here you can host XR applications to be used on the campus. When an application is hosted on-premise for external use (i.e., it needs to be accessed off-campus), you generally have to request a firewall change.

2.1.2.3 Elsewhere

Outside of the campus, we encounter a heterogeneous landscape. XR HMDs gain network connectivity through public WiFi or mobile 4G/5G telecommunication networks. (As already mentioned above, the current generation of XR HMDs do not support 4G/5G connectivity, and we need to use WiFi hotspots or MiFi routers to make use of these mobile networks.) As network traffic is routed over the public internet, it is a low-trust environment. Protocols like uPnP are not supported in these environments, so only applications that use IP-based routing will work in this context. Finally, we can host XR applications on edge servers or in the (public) cloud that are easily reachable because we tend to have more control over firewalls and other restrictions in these settings.[7]

[6] At the same time, it could be possible with eduroam and iotroam to add devices on different locations on the campus to the same network, depending on the network architecture of the campus.

[7] In the context of this report, we define the cloud as a central datacentre where an application can be hosted on servers. The cloud could be one of the public cloud providers (e.g., AWS, Azure, GCP), but also a privately-owned datacentre or community cloud, such as the cloud services offered by SURF. In contrast, the edge refers to servers (and other compute devices) that are located nearby end-users. Although there is a grey area in these definitions (for example, should an on-campus datacentre be regarded as edge or cloud?), the key takeaway is that it is important to consider the location of where an application is running in relation to the end-users.

3 5G and WiFi 6

The new wireless technologies, 5G and WiFi 6, help to scale XR in different ways. With 5G, it will become possible to use new types of XR applications, in particular remote rendering and multiplayer applications, in a variety of (outdoor) locations, and with WiFi 6, we can scale up the number and usage of XR HMDs indoors.

5G is designed to support three usage scenarios: enhanced massive broadband (eMBB), ultra-reliable low latency communications (URLLC), and massive machine type communications (mMTC). XR technologies can benefit from all three of these scenarios. Streaming to XR HMDs is improved with eMBB, URLLC makes interactive and multiplayer applications possible, and with mMTC a huge number of sensors can be used as inputs for XR applications. According to the 5G specification, 5G should have a theoretical maximum bandwidth of 10 Gbps compared to 150 Mbps for 4G, and a maximum latency of 4ms versus 20ms for 4G.[8] Although these limits assume perfect circumstances, they do indicate that outdoor mobile connectivity will approach the same quality as our indoor networks and will make it possible to use XR applications outdoors.

The improvements that WiFi 6 brings will help scale the indoor usage of XR HMDs. The theoretical maximum bandwidth of WiFi 6 is 9.6 Gbps compared to 6.9 Gbps for WiFi 5, which admittedly was already rather high. More importantly, by using orthogonal frequency-division multiple access (OFDMA) and multiple user-multiple-input and multiple-output (MU-MIMO), WiFi 6 is better able to increase the throughput for several HMDs in the same area.[9] These improvements result in a (theoretical) latency reduction of 75% compared to WiFi 5.[10] In other words, WiFi 6 will play an important role in scaling up the number of XR HMDs that can be serviced by one WiFi access point on the campus.

In summary, both the 5G and WiFi 6 specifications indicate that the current developments in wireless technologies will help in scaling up XR soon.

[8] https://www.itu.int/md/R15-SG05-C-0040/en
[9] https://arxiv.org/pdf/1501.01496.pdf
[10] https://www.zdnet.com/home-and-office/networking/next-generation-802-11ax-wi-fi-dense-fast-delayed/

4 State of the art: XR hardware and software

As of this writing, there is a variety of XR hardware and software available that could suit this experiment. To make an informed decision, we evaluated a few options in both the hardware and (XR streaming) software categories.

First, we wanted to use off-the-shelf solutions wherever possible. A lot of this would be achievable with custom hardware and custom software. However, the goal was for the findings of this report to be easily repeatable and implementable across academic institutions.

Unfortunately, as of the start of this experiment, there were no mainstream commercially available XR HMDs with integrated 5G antennas available. This means that all solutions would require a 5G to WiFi (MiFi) modem between the 5G mast and the HMD. In our experiments, we made sure to limit the number of HMDs connected to a single modem to ensure no bandwidth limitations; however, additional latency must be accounted for due to the addition of a WiFi router.

XR HMDs can broadly be divided into two sets: Standalone and Tethered. For this experiment, due to the “in the field” nature, we decided to focus our attention on standalone HMDs only. We evaluated the Meta Quest 2 and Pico 4. We did not include the Vive Focus 3, as it was not yet released at the start of the experiment. We focused only on hardware, price and technical capabilities, and disregarded factors such as privacy policies. Both of the mentioned HMDs support WiFi 6, with the Meta Quest 2 having a broader range of supported bands. This fact, combined with the lower price, made us decide to use the Meta Quest 2 for this experiment.

In terms of streaming software, Meta has its own streaming protocol, called AirLink, available for wirelessly streaming XR content. Other popular options are Virtual Desktop and ALVR. We also wanted to include nVidia’s CloudXR platform, but unfortunately, we were not granted access in time. We tested AirLink, Virtual Desktop and ALVR in this experiment.

5 5G test project

The Do IoT Fieldlab has created a testbed environment for testing 5G use cases. The field lab consists of two locations: the Green Village at the TU Delft and the Unmanned Valley in Katwijk. The fieldlab offers three advantages for testing future use cases of 5G. First, the fieldlab has a license to use all three of the 5G frequencies. The EU has designated three frequencies for the use of 5G: 700 MHz, 3.5 GHz, and 26 GHz. At the time of conducting our project, the 3.5 GHz and 26 GHz were not yet available in the Netherlands.[11] The 3.5 GHz frequency is particularly valuable to work with for the improvements it brings in bandwidth and latency compared to the 4G spectra, while offering coverage than 5G's 26 GHz frequency. Second, the fieldlab provides the ability to deploy applications to the edge of the 5G network. These 5G edge deployments are part of the 5G specification and are meant for use cases that require fast and reliable connectivity.[12] Third, the fieldlab offers expert support, consultancy and a controlled environment for testing and troubleshooting 5G use cases.

As already mentioned in the section “5G and WiFi 6”, the improved capabilities of 5G create interesting use cases for XR. In our tests we decided to focus on enhanced massive broadband (eMBB) and ultra-reliable low latency communications (URLLC) because these usage scenarios aligned well with existing XR use cases in education that would benefit from scaling. For the eMBB scenario, we selected a remote rendering test. As the compute capabilities of XR HMDs are constrained, certain simulations will have to be done on external hardware, e.g., servers in the public cloud or 5G edge. For the URLLC scenario, we selected a multiplayer test, as this type of application benefits the most from low latency connectivity. In addition to 5G, we also took the broader challenge of scaling XR into consideration in our test cases.

5.1 Use case 1: Remote rendering

The first use case we tested was remote rendering. Remote rendering could help with one of the largest pitfalls of modern XR HMDs, namely their limited rendering capabilities. A standalone XR HMD runs on a mobile processor, which at the time of writing is not capable of rendering highly realistic XR environments at 90 Frames per second. With this use case, we explored the possibility of rendering remote content from both a PC connected to the 5G edge and a Virtual Machine (VM) in the public cloud and streaming their content to an XR HMD via 5G.

5.1.1 Test criteria

To evaluate our remote rendering tests, we used the following criteria:

- Connectivity: the HMD can connect to the streaming server.
- Stable connection: during streaming a user does not experience network interruptions.
- Responsiveness: XR experiences are very personal and what is responsive depends heavily on the use case. As a rule, delays should be less than 100ms in the context of research and education. We used the latency metrics provided by the streaming software, Virtual Desktop, to gauge responsiveness, as well as subjective experience.
- Smooth movement: movement should follow the user's movement without too many shocks.

[11] https://www.rijksoverheid.nl/onderwerpen/telecommunicatie/invoering-5g/veiling-frequenties-voor-5g

[12] It remains to be seen whether the telecommunication companies will add the functionality of edge deployments to their public 5G networks. It could be that this functionality will be limited to private 5G networks.

- No lag spikes: sudden moments or a drop in the response time can ruin/disturb the experience.

5.1.2 Test setup

In this section, we discuss the test setups we used in our remote rendering test from a development and infrastructure perspective.

5.1.2.1 Development

For our tests we needed to figure out three components: 1) how to connect HMDs to a rendering PC/VM, 2) how to connect an XR HMD to the 5G network, and 3) select an application that is suitable for streaming XR.

There are multiple tools for connecting an HMD to a PC/VM, such as AirLink, Virtual Desktop, and ALVR. These tools are designed for personal use in a home environment and, at this moment in time, there are no good solutions for streaming XR applications to multiple XR HMDs. Looking at these tools, Virtual Desktop proved to be the most versatile and capable of connecting over IP. Virtual Desktop also provided real-time benchmarks of the quality of the connection, which was useful for our quantitative measurements.

The second component, bringing 5G connectivity to XR HMDs, was much easier to solve. For our tests, MCS supplied and helped set up Advantech ICR-4453 routers, which is a 5G Modem and WiFi router in one. The router connected to the 5G antenna at the testing sites and converted this signal to a WiFi network to which the XR HMDs could connect.

Lastly, for the XR application we selected “The Lab” from Valve as the main testing experience. “The Lab” is a VR game that is accessible for everyone and allows for a more repeatable experiment.

5.1.2.2 Infrastructure

Before discussing the infrastructure setups for remote rendering from the 5G edge and from the public cloud, we would like to point out that these setups deviate from the ideal scenario in three ways:

1. We introduced a router to connect the XR HMDs to the 5G network because the current generation of HMDs do not support 5G directly.
2. For testing purposes, we opted to use workstation from our lab to render content on the edge of the 5G network. This way we were sure that the quality of the VR output would be good because the PC’s output has been validated in the lab. In the ideal scenario, we would have deployed our remote rendering application on a 5G edge server but the servers in the fieldlab did not have GPUs.
3. When streaming from a VM in the public cloud, we needed to use an additional workstation to create a display output for the XR HMDs. For similar reasons, the PC streaming content from the 5G edge needed a monitor. Ideally, the additional PC and monitor would not be necessary for streaming the XR content.

5.1.2.2.1 Setup for remote rendering from the 5G edge

For this test we placed a workstation near the 5G edge server located in the Green Village of the TU Delft. We used our own workstations as this would allow us to compare the performance using the HMD and the PC through a WiFi 6 modem as well. The HMD was connected to a local

WiFi router with a built-in 5G modem, this router connected to the 5G test network at the Green Village.

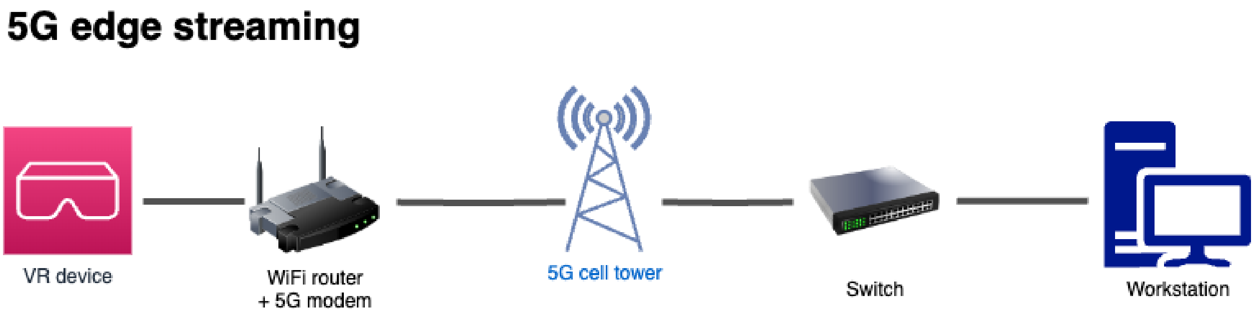

Figure 1. 5G edge streaming single device setup, Source: Tim Kok

In the image below, you can see the setup we used for streaming to multiple HMDs from the 5G edge. This setup was similar to the single HMD edge streaming setup but had multiple parallel connections using the same cell tower and switch to connect to multiple workstations. The image shows that this setup is not suitable for scaling up because each XR user needs a separate workstation for rendering content. In future setups, we would like to change the setup into a more scalable solution, like VMs located in the edge or cloud.

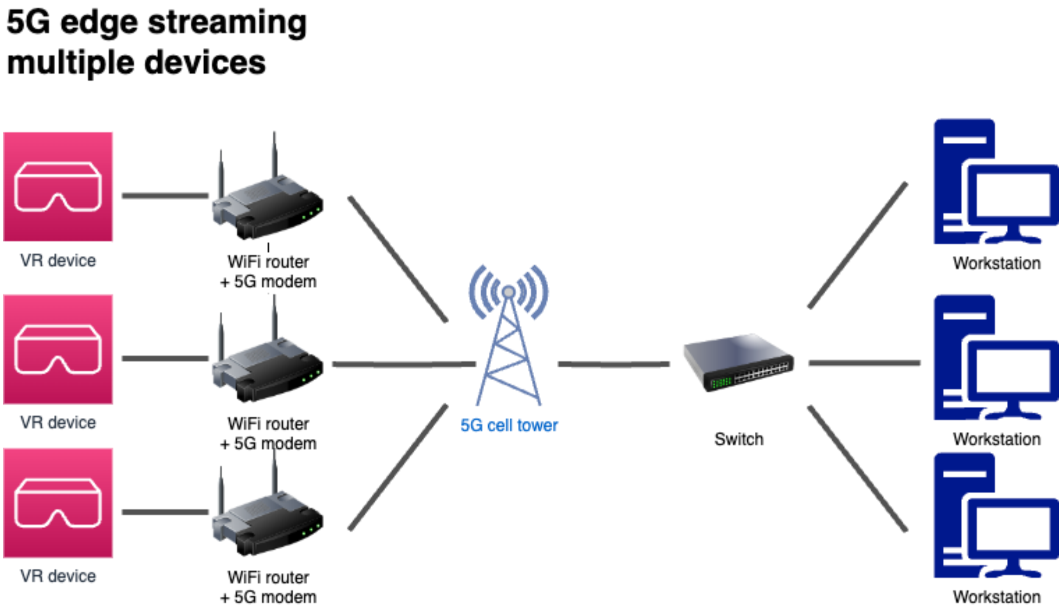

Figure 2. Multi-device streaming setup using the edge, Source: Tim Kok

5.1.2.2.2 Setup for remote rendering from the public cloud
To test the possibility of streaming from the public cloud, we deployed a Windows VM with GPUs on AWS and installed Virtual Desktop on the VM. This setup did not work out of the box because Virtual Desktop needs a display output from the PC to be able to determine how big the virtual screen is and how to stream the content to an HMD. Attaching a physical monitor is not an option for a VM, but remote desktop tooling also creates a (virtual) display output. We first tried creating a display output with RDP, but this resulted in a green screen on the HMD because

RDP selects a screen driver that is incompatible with Virtual Desktop.[13] After testing with Moonlight and TightVNC, we found that Parsec worked best. Parsec needs to be installed on the VM and a workstation before being used to stream content from the VM to the HMD. After solving the display output issue, we managed to connect the XR HMD to the VM and play games in “The Lab”. We did experience some large latency spikes, but most of the time the XR experience was smooth and usable. To give a conclusive answer on what VM setup in the cloud would be best further testing is needed.

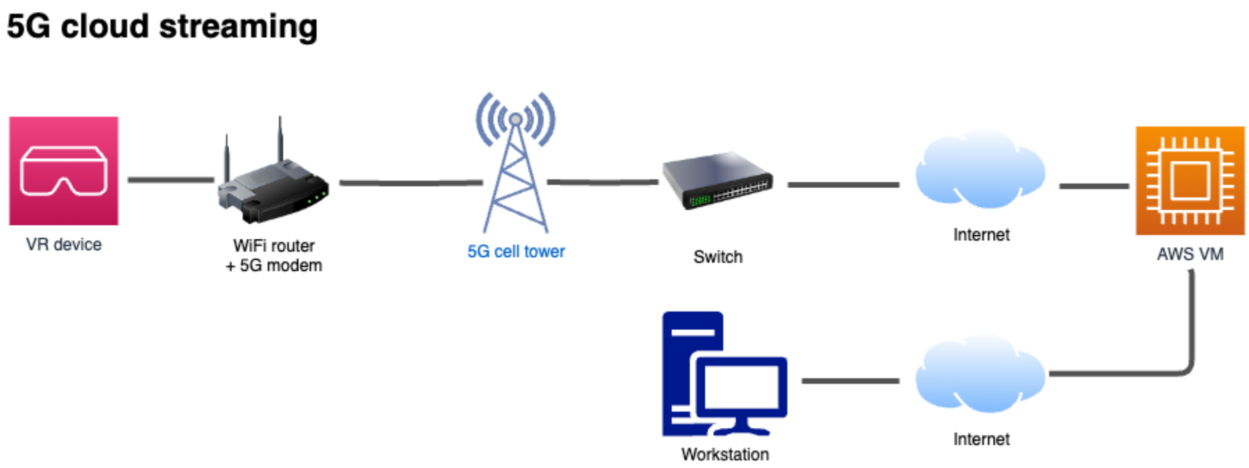

Figure 3. Cloud streaming setup using the cloud, Source: Tim Kok

As the image below shows, scaling up this setup requires setting up a VM and workstation per HMD.

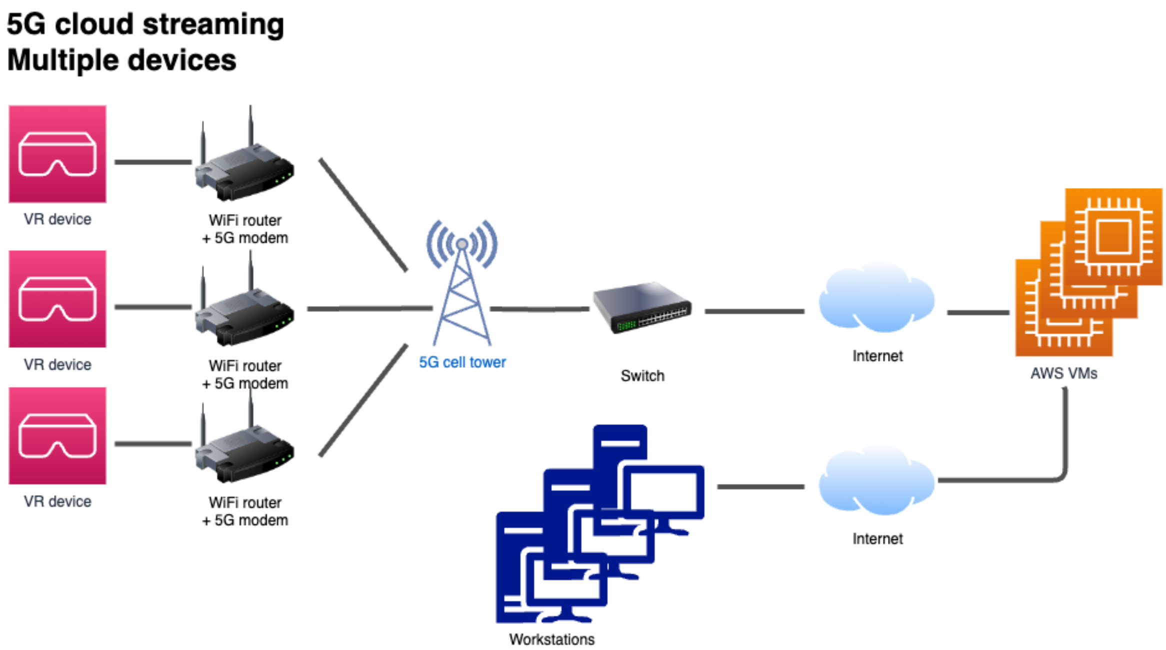

Figure 4. Multi-device cloud streaming setup, Source: Tim Kok

[13] This issue was also encountered by other users, see for instance https://www.yellow-bricks.com/2020/01/02/seeing-green-only-on-your-hmd-when-using-alvr-to-stream-an-app/.

5.1.3 Results

The PCs on the edge were sufficient for most medium speed experiences. "The Lab" from Valve provided a wide range of interactions. The user could move objects, shoot a bow, and do all interactions that are available in "The Lab". However, there was a small amount of lag, the hand would move behind your real hand in a smooth motion. This would lead to a discrepancy between visual input and sensory input of where your real hand is. For example, when waving your hand, your virtual hand would not do the full range of motion of the wave but only 80% of the movement. This is not a problem for most experiences, but in case the experience needs very low reaction times and precision this could be a problem, e.g., in a F1 racing simulator.

Another difficulty we encountered is that each HMD required a separate PC or VM. This is a limitation of Virtual Desktop, which connects one HMD to one PC. Possibly, this limitation can be solved with software designed to stream to several HMDs at once, but this software was not available at the time of testing.

The setup is difficult to scale. When connecting more than 4 HMDs to the 5G router we experienced a noticeable drop in performance, e.g., lower frame rates and more latency. Which is most likely a limitation of the 5G router. A better router could be used to add even more HMDs to the setup. When XR HMDs get built-in 5G modems this problem of scaling would be solved.

5.1.4 Discussion

5.1.4.1 5G connectivity

A large limiting factor that contributed to a lot of extra overhead was the absence of a built-in 5G modem on the HMDs as this added additional steps and expensive hardware to the whole setup. Hopefully, in the future, HMDs will have an integrated 5G modem or an affordable, small, 5G modem needs to be used for each HMD (or larger routers for multiple HMDs).

5.1.4.2 Virtual machines

Virtual machines are the most flexible, accessible, and (depending on usage patterns) affordable solution to solve the backend pitfalls of the current setup.[14] The VMs can be used to reduce the number of setups near the edge and/or could use the native hardware of the edge to compute the XR experiences. The exact setup of the VMs is use case dependent, but a portal could make the whole process of using VMs easier and more viable for smaller-scale institutions by automating the configuration, deployment, and connection to a specific VM.

At the same time, the limitations of using VMs, especially in the public cloud need to be tested further. During some of our tests, streaming from a VM on AWS worked well, but at other times this resulted in large latency spikes, ranging from 40ms to 400ms, which also affected the user experience. This is too much latency for XR use cases and needs to be researched further for VMs to be considered a stable option. It is unclear whether these spikes were caused by the

[14] Whether a VM is indeed cheaper (or better) than alternative options, such as purchasing workstations or on-premise servers with GPUs, needs to be determined per use case. As a rule, small, short-lived applications or applications with irregular usage spikes are well-suited for the cloud, and as the resource demands become larger, more consistent, and more permanent, on-premise solutions tend to be more advantageous. Other considerations could be the expected availability of the resources, e.g., GPUs are at times in short supply in (public) clouds and may need to be reserved ahead of time, making it a less flexible setup.

limited resources of the VM (we selected one of the smaller VM sizes on AWS), the geographic location of the VM (Frankfurt), or other (networking) factors.[15]

5.1.4.3 Setup

Creating a test setup for 4 HMDs took a lot of time. For each user, the team needed to install a workstation, and later deploy and configure a VM, to be able to offload the rendering. For 4 HMDs it took half a day to properly set up. With, for example, 25 HMDs this would cost too much time. A tool that would set up VMs for the user would help solve this issue and might reduce the overall costs of the setup as the user would only pay for the setup when in use.[16] Taking these benefits into consideration as well as the possibility to simulate situations that require a lot more computational power, could make VMs a viable option.

5.2 Use case 2: Multiplayer

Most XR applications are single-player and miss the aspect of interaction between players. However, at the XR Zone, we developed several multiplayer experiences, and we believe future XR applications, inside and outside of the education system, will be made with collaboration in mind (in effect multiplayer). Therefore, the second use case tests the low latency improvements 5G brings and uses a multiplayer XR experience to explore if 5G can provide the low latency XR multiplayer applications require.

5.2.1 Test criteria

To evaluate our multiplayer tests, we used the following criteria:

- Multiple HMDs need to be able to connect to the same game instance hosted on the Application management server (AMS).
- Stable connections.
- Real-time updates and good enough responsiveness for XR in education and research.

5.2.2 Test setup

In this section, we discuss the test setups we used in our multiplayer tests from a development and infrastructure perspective.

5.2.2.1 Development

For the multiplayer tests, the XR Zone supplied a standard Unreal Engine template with physics actors to test if the actors would be replicated properly and responsively across the HMDs. In our setup, we need to distinguish between the functionality of the “Application Management Server” (AMS) and the “Game Host” functionality. The AMS is used to start a game and add HMDs to specific games. The Game Host keeps track of the state of the game and communicates this to all the players, e.g., what all players are doing and the location of the various objects and actors. In our test setup, the AMS was running on a separate server accessible from the internet and one of the HMDs had the role of Game Host.

The development was mainly focused on making sure all the HMDs would be able to connect to the HMD hosting the XR game. To connect the HMDs together, they need to be able to find each

[15] Remote rendering in the XR Zone lab and from a 5G edge server worked well, so these spikes were somehow caused by the VM setup in the public cloud.

[16] Cost savings in the cloud depend on the context and use case. See footnote 14.

other. This can be done in two ways:

1. Locally, look within the local network for the server or players and establish connections.
2. Connect by IP, use an IP to connect an XR HMD to another by sending the AMS server the IP of the client.

Unreal Engine supplies a multiplayer connection system, but not a connection over the internet for small-scale projects. Therefore, the XR Zone developed their own web server and portal, the AMS, to send the IP from one XR HMD to the other, allowing them to connect by using a short passcode. The AMS could also overwrite the IP of an HMD to deal with proxies and Network Address Translation (NAT).

5.2.2.2 Infrastructure

We could have taken two kinds of approaches to make the multiplayer of multiple HMDs work:

1. Placing the Game Host on the HMD and allowing other HMDs to connect to this “game server HMD”.
2. Hosting the game on a central server and connecting all the HMDs to this server, which would send all the instructions and game states to the HMDs.

Eventually, we decided to place the Game Host on the HMD (approach 1), as Unreal Engine can do this easily and cheaply, and with this approach there was no need for a separate VM to host the game. Moreover, if the first approach would work well, it is reasonable to expect that the second approach would yield similar (if not better) results due to the increase in compute resources of using a dedicated VM.

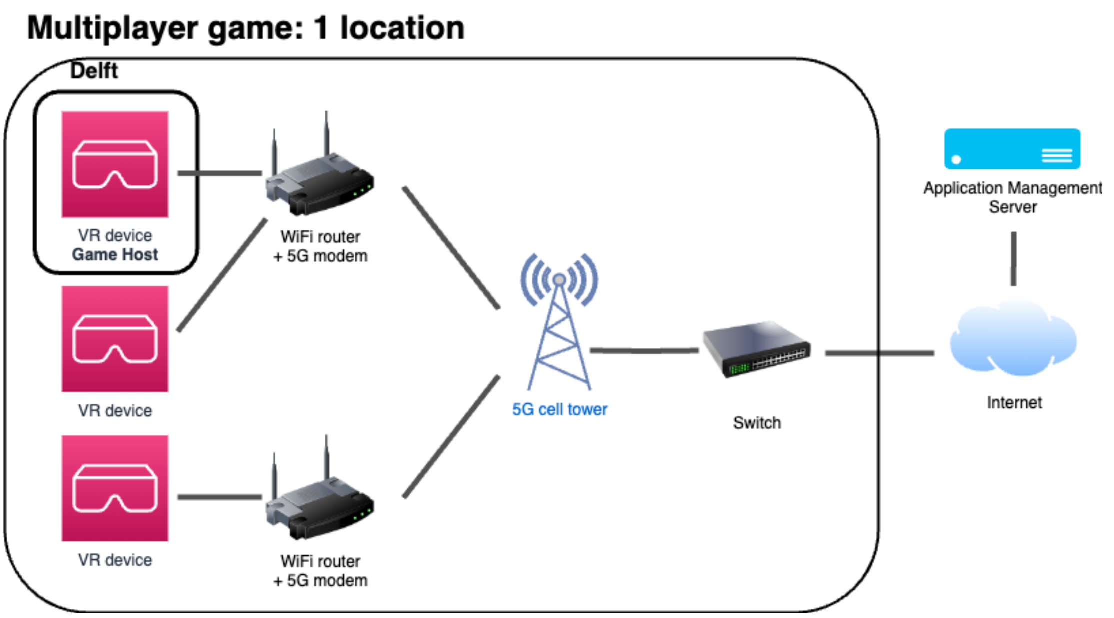

Figure 5. Multiplayer, one location, over 5G setup, Source: Tim Kok

In the first test setup, multiple HMDs were connected via 5G to the same Game Host in one location. The AMS made sure the HMDs could find each other by IP. As shown in the image above, not all HMDs were connected to the same router as this would give us more insight into the connectivity options within a 5G network. For the HMD connected to a different router than the Game Host, we needed to use the router's DMZ feature[17] to establish communication with the Game Host.

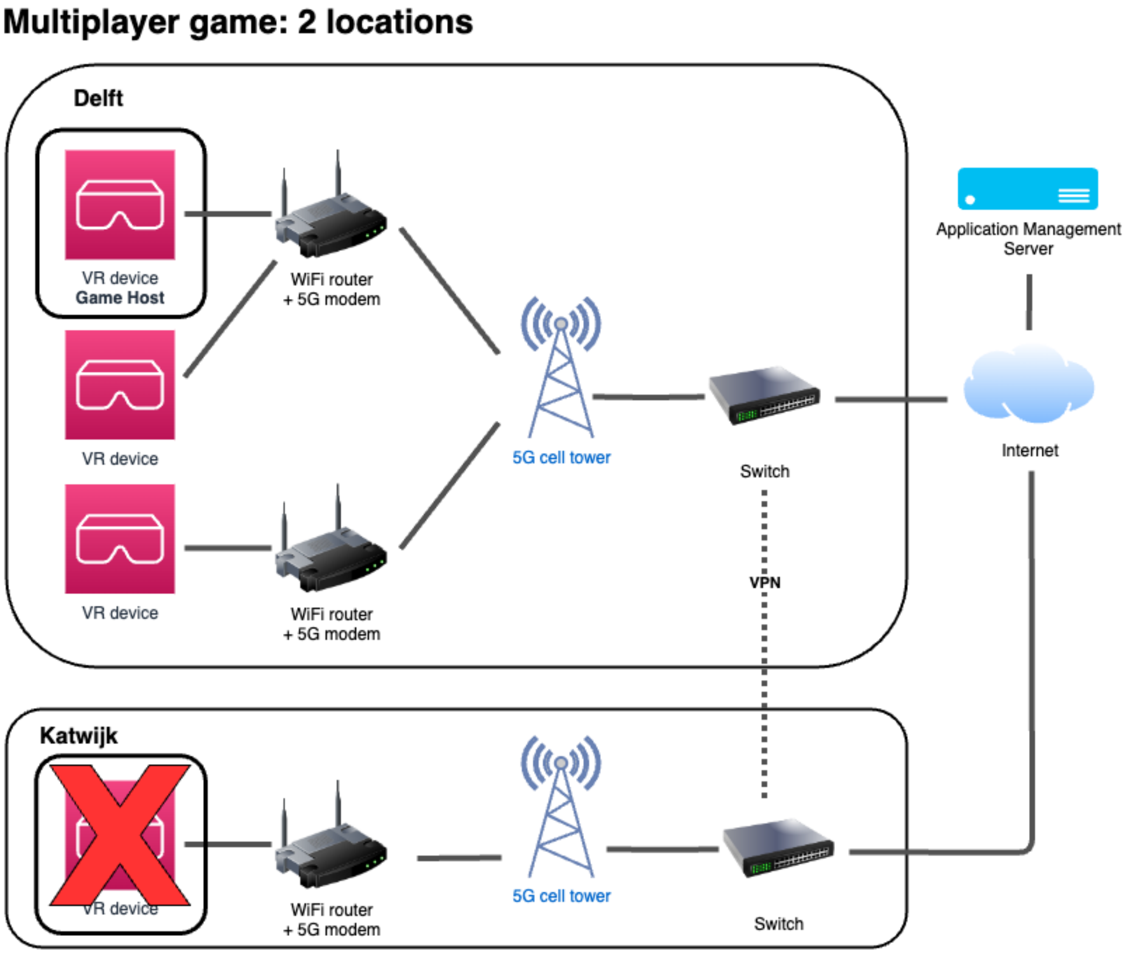

Figure 6. Multi-location multiplayer over 5G setup, Source: Tim Kok

In the second setup, we tried to connect HMDs from two locations to the Game Host present at one of the locations. Unfortunately, this setup did not work as the second location (Katwijk) was not able to connect to the Game Host running on the local HMD at the TU Delft. The reason for this was that the (private) 5G networks in TU Delft and Katwijk could not route traffic between each other, which is also likely to happen in public 5G networks.

5.2.3 Results

The main difficulty was for the XR HMDs to find each other on the network. The 5G network was capable enough to send the limited amount of data required for a multiplayer game between the HMDs, but connecting to the Game Host was difficult when the Game Host was running on an HMD.

We set up a system on the AMS that sends the IPs of the Game Host HMD to the other HMDs and would have the HMD ask for a connection directly with a passcode for safety. However, this system is not sufficient when HMDs are located in different private networks. The most

[17] https://en.wikipedia.org/wiki/DMZ_(computing)

straightforward way to solve this issue is to have a separate game hosting server that can be reached by all HMDs.[18]

In the future, a good solution might be to combine the VM portal, proposed in the previous section for deploying rendering VMs, with a feature that deploys a game hosting VM that will be reachable by all the HMDs (i.e., approach 2). This setup will also shift the workload from the HMD to the VM, reducing the load on the already limited processing power of the HMD.

5.2.4 Discussion

In contrast to the remote rendering use case, multiplayer use cases require less bandwidth from the network and would only benefit from the reduced latency of 5G. This raises the question whether the 5G improvements are a necessity for multiplayer use cases or not. In other words, would the multiplayer tests work on a 4G network? Although testing multiplayer games with 4G was outside of the scope of our tests it could be an interesting topic for future research.

With this setup, it is very important to have developers that understand multiplayer applications and the complexity that a multiplayer setup brings. The offloading of the server and sharing the data between HMDs with a server requires careful programming and a well-thought-out program. In addition, data transfer between all HMDs needs to be optimized to be able to use the experiences even with unstable connections. The tests with 5G had a stable environment that had almost no jitter, but when using a public 5G network this most likely will not always be as stable.

[18] Another option would have been to use VPN tunnels or other exotic network configurations, but these types of solutions do not work in all contexts and add a lot of complexity to the setup.

6 Conclusions

In this final section, we will highlight some of the overall findings from our tests and draw some broader lessons.

6.1 XR and 5G

Based on our two test cases, we can conclude that 5G networks likely provide sufficient bandwidth and low latency for XR applications in education and research. We reached a maximum bandwidth of 190mbit/s. The latency between the HMD and the 5G edge server, as reported by Virtual Desktop, stayed between 30ms and 80ms. The latency between the HMD and a VM in the public cloud (located in the Frankfurt region), mostly stayed around 70ms – 80ms with at times short spikes of 100ms – 120ms.

Nevertheless, we expect the actual usage of 5G in XR to be limited, at least in the short to medium term. The adoption of 5G in XR will depend on changes in XR hardware and software, as well as broader societal changes. First, XR HMDs will need to include a 5G modem and support outdoor usage. Movement and environment tracking will need to work in broad daylight. Currently, the Meta Quest 2 controllers are tracked via infrared sensors (which are unreliable in broad daylight), and movements are tracked via cameras that are too sensitive to be used outside. Dedicated chips to optimize video decoding, e.g., Apple's A17 Pro chip,[19] and sensor processing, e.g., Apple's R1 chip,[20] might also improve the capabilities of the XR HMDs in the future and help with adopting streaming as a viable way of using XR HMDs.

Another potential push for 5G would be if people switch from WiFi to 5G for their internet connectivity at home. A large group of users may not need the fast internet WiFi 6 in combination with fibre optics provides, but just need internet access. For these homes, a stable 5G connection could be sufficient and this could drive 5G support for consumer devices in general, such as XR HMDs.

Augmented reality (AR) applications are another potential factor that could positively affect the adoption rate of 5G. AR applications are often based on real life locations and would benefit from a stable, fast, and low latency connection to a (rendering or multiplayer) server. At this moment in time, AR HMDs and applications are not integrated into our daily lives and need to improve to see more adopters latch on. However, when AR hardware and software improve, we could see a large push for 5G connectivity for AR.

As it stands, most current XR HMDs and XR applications are developed for indoor use, where WiFi is generally available. Therefore, in the context of the campus and lab, we expect that the transition from WiFi 5 to WiFi 6 will have a bigger impact than 5G on XR, at least in the near future.

[19] https://www.apple.com/newsroom/2023/09/apple-unveils-iphone-15-pro-and-iphone-15-pro-max/
[20] https://www.apple.com/newsroom/2023/06/introducing-apple-vision-pro/

6.2 Software limitations

Another general finding from our tests is that the current state of XR software and tooling is not quite ready to support usage outside of a home environment. Indeed, XR software and tooling are probably the biggest bottlenecks in scaling XR for education and research. A few changes that would help to use XR at scale are:

- Software that can handle routing based on IP, domain names, and NAT, as an alternative to uPnP.
- Software that supports more advanced authorization and authentication scenarios, such as an integration with a custom identity provider to avoid the use of shared accounts.[21]
- Streaming software that can host multiple applications on 1 VM and/or a cluster of VMs and can stream to multiple HMDs simultaneously (preferably using a native streaming integration, such as AirLink).
- Management tooling for a fleet of HMDs that, for example, supports setting policies on which HMD can (or should) access which application.

6.3 Edge deployments

Our tests also showed that the locality of the XR applications can be an important factor when designing an XR setup. Remote rendering on workstations at the 5G edge was significantly better in terms of latency, stability, and subjective experience. Although our tests could have been more extensive, for instance, we only tested one type of VM in one AWS Region,[22] the results are in line with expectations. Therefore, we recommend deploying XR applications for remote rendering and multiplayer on edge servers on campus or in nearby data centres when performance is critical.

6.4 Privacy and security

Privacy and security were outside of the scope of this project, but they are important aspects when scaling up the use of XR in education and research, so we briefly mention it here. As soon as we leave the trusted environment of an XR lab, basic security standards should be in place, e.g., encryption and robust authentication. Similarly, privacy concerns regarding the analysis and storage of sensitive information need to be addressed before XR applications can truly be used at scale by students and research participants.

6.5 XR streaming

Regarding XR streaming, we would like to make two remarks. First, we expect XR streaming HMDs to take a larger part of the total XR HMD market in the future. These HMDs would be slim and only capable of streaming XR experiences from other locations (local devices or cloud). We can already see how small these HMDs could become with products like the Bigscreen Beyond.

Second, although our report considered mobile network technologies and (public) cloud, we want to highlight that the first step when developing an XR application should be optimizing the

---

[21] Virtual Desktop uses a Meta account to give an HMD user access to a computer, but within larger organizations a Meta account is not necessarily unique or linked to a specific user.

[22] "Amazon cloud computing resources are hosted in multiple locations world-wide. [...] Each *AWS Region* is a separate geographic area" (https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.RegionsAndAvailabilityZones.html). For our tests, we deployed a VM in the Frankfurt region.

XR application itself. In most cases, an application should be able to run on the HMD itself. This approach reduces the cost and complexity of a project. Of course, if running on the HMD is not feasible, streaming is a valid option. Moreover, in the case of multiplayer games at scale, using cloud infrastructure is almost unavoidable.

6.6 XR multiplayer

The main finding from the multiplayer test case was that communication between HMDs is a problem in the context of different private networks because of IPs that are not routable. A common solution to this problem is to use a central server to which the individual HMDs connect as clients and that takes care of communication. The main obstacle here for XR developers is that there is no standard solution in the XR ecosystem to deploy a multiplayer application on a server in an easy way. Currently, this process requires knowledge of system administration and software deployment. When the services and tooling for deploying and hosting XR multiplayer applications improve, this will likely boost the number of multiplayer applications for XR.

6.7 Potential next steps

There are several avenues that could be explored further in this area. The first would be tests looking more closely at WiFi and XR. For instance, comparing the performance of WiFi 5 to WiFi 6, and testing the maximum number of XR HMDs one WiFi access point can support. Another avenue would be exploring and/or developing alternative technologies for scaling XR applications in the cloud, such as containers and WebRTC for XR. Finally, we think it would be valuable to develop new use cases in the domain of education and research for XR streaming, such as simulations for fluid dynamics and black holes.

6.8 Final thoughts

How 5G and 6G mobile networks will shape the XR landscape in the long term remains to be seen. At the same time, the opportunities that 5G and (soon) 6G provide are seemingly endless, and we see a future in which XR HMDs will be comfortable, small, and connected because of technologies like 5G, edge and cloud. In the meantime, there are plenty of challenges that need to be solved to start scaling XR in education and research.

7 Credits

We would like to thank the following people and teams for their support:

Pim Zwager (MCS)
Floris Drijver (TNO)
Joep van der Weijden (The Green Village, TU Delft)
Marijn Leeuwenberg (The Green Village, TU Delft)
Yosua Pranata Andoko (XR Zone, TU Delft)
Sharif Bayoumi (XR Zone, TU Delft)
Paul Melis (SURF)
SURF Cumulus

Publication | Januari 2024

SURF
Kantoren Hoog Overborch
Moreelsepark 48
3511 EP Utrecht
088 - 787 30 00