



BILBOKO
INGENIARITZA
ESKOLA
ESCUELA
DE INGENIERÍA
DE BILBAO

**MÁSTER UNIVERSITARIO EN
INGENIERÍA DE TELECOMUNICACIÓN**

TRABAJO FIN DE MÁSTER

***LARGE SCALE IMMERSIVE HOLOGRAMS
WITH MICROSOFT HOLOLENS***

Alumno/Alumna *Lu, Lu, Andrés*

Director/Directora *Hajek, Jeremy R.*

Departamento **School of Applied Technology**

Curso académico **2018/2019**

Chicago, Agosto 2019

I. Index

I.	Index.....	2
II.	Abstract.....	4
III.	Introduction.....	5
IV.	Background	8
V.	State of the Art	12
	State of the Art Conclusions	16
VI.	Objectives.....	18
	Feasibility of LIDAR-scanned objects to act as LSIH.....	18
	Point-cloud rendering technique analysis with FARO SCENE	19
	Mesh objects polishing and adaptation to Unity3D with Meshlab.....	19
	Unity3D setup and performance tests with HoloLens Device/Emulator	19
	Spatial Mapping concepts and tests	20
	Future technologies/devices for LSIH.....	21
VII.	Methodology	22
	Feasibility studio and analysis of LIDAR scanned files to act as LSIH.....	22
	Point-cloud rendering techniques available in FARO SCENE	25
	Importing the scans	26
	Processing the scans.....	26
	Registering the scans and creating mesh objects from the 3D view	28
	Mesh objects polishing and adaptation to Unity3D using MeshLab.....	32
	Smoothing and remeshing techniques.....	34

Unity3D setup and performance tests with HoloLens Device/Emulator	38
Requirements to deploy an Unity3D project to HoloLens	38
Configuring the project settings	39
Deploying the project to Visual Studio	41
Spatial Mapping concepts and tests	44
Description of future technologies for LSIH	46
VIII. Results	48
LSIH performance test results	48
Spatial Mapping test results	52
Future technologies and devices for LSIH	58
Microsoft HoloLens v2	58
Cloud rendering/computing technologies	62
First option: external GPU unit	63
Second option: cloud computing/rendering solutions	66
Cloud solutions in the market	71
Google Cloud Rendering (Zync Render)	72
Amazon Sumerian:	73
GarageFarm / Xesktop:	74
Fox Renderfarm:	75
Rendershot:	75
IX. Conclusions	77
X. Bibliography	79

II. Abstract

This project focuses on the premise of researching whether or not a Large Scale Immersive Hologram based off a 3D model obtained via LIDAR scanner technology is viable to be deployed, currently or in future iterations, in Microsoft's HoloLens device. For that purpose, we will be addressing how LIDAR technology works, what can be obtained from the 3D models it generates and how we can polish and optimize the resulting tridimensional mesh objects. Later on, we will implement these objects in a development environment compatible with the HoloLens device, Unity3D, and run a performance test to see how suitable and realistic the user experience results. Furthermore, we will research how near-future technologies can greatly help to enhance this experiences through future iterations of this same device.

Keywords: LSIH, AR, MR, HoloLens, Holograms, LIDAR, 3D

III. Introduction

Augmented Reality technology is used to experience interactions with the real world by expanding it with the use of objects that could be rendered by a computational system. Opposed to Virtual Reality, this approach mixes both realities: virtual and real, into a single experience. The way in which reality can be “augmented” is huge, and almost limitless. The perceptual information that can be generated to augment reality can be based upon plenty of our sensory modalities: visual, auditory, tactile, olfactory... This information is then overlaid upon the real sensorial information the subject is perceiving, thus mixing both realities. This difference is key in understanding how Augmented Reality is different from Virtual Reality. AR does not isolate the subject in another reality. Instead, it enhances and works over the existing reality. This is precisely what diverges from VR and the main point why AR applications can have a very different focus.

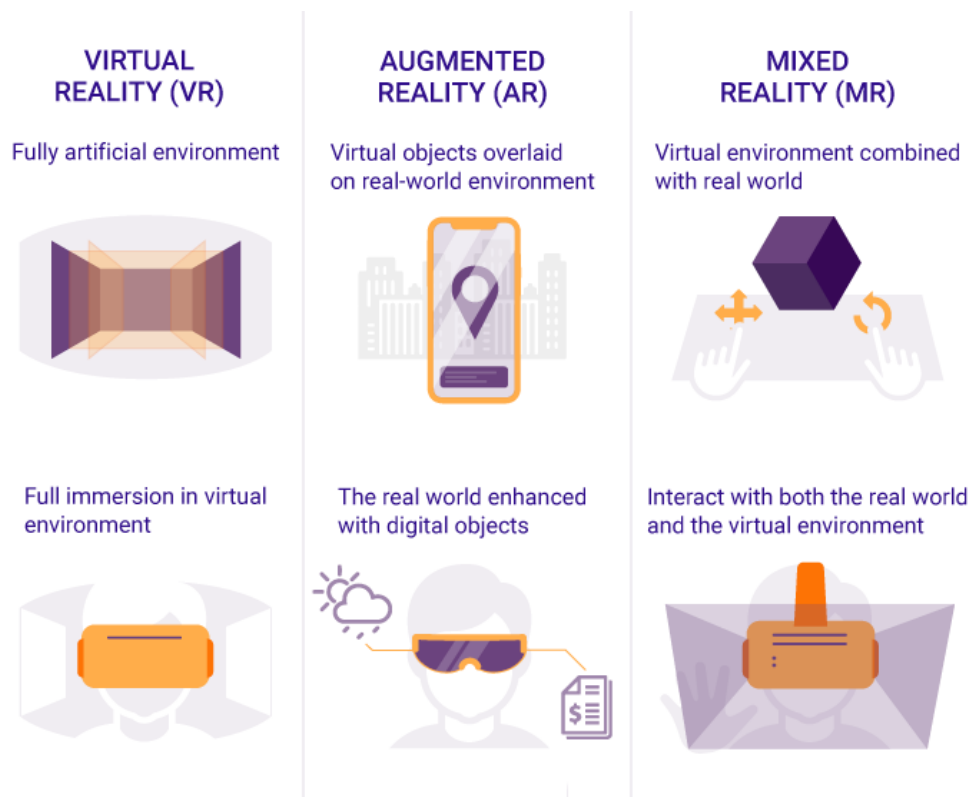


Figure 1: Difference between VR, AR and MR

AR can help us enhance the way we perform certain tasks, or even enable us to do things that were not possible before. This is precisely the focus of this particular review, to seek out in which ways could AR technologies, in particular Microsoft's HoloLens devices, to improve quality of life, performance and production in already existing tasks and how this technology can help us do things never imagined before.

This particular project could be categorized as both AR or MR (Mixed/Merged Reality) depending on the approach used. The main difference between AR and MR would be the interactivity of the "virtual" environment or appliances with the real world. Meaning that if we were to project a hologram through a device such as Microsoft HoloLens to be a static hologram, always restricted to the same size or fixed location, it would classify as AR. However, if we made that particular hologram/object vary with the surrounding environment, for example, by changing its size or scale based on how extensive the floor we are looking at is, that would classify as MR since we are using real-world data to interact with our virtual object: we are creating an interaction between the real and virtual environments.

Although for this particular project VR could also be used, AR has several advantages that could fulfill the purpose of displaying holograms in a better way than them being in a fully virtual environment. Some of them are:

- Overlaying virtual objects to the real world gives a greater sense of perspective and sizing than working in a fully virtual environment.

- AR is one step ahead when speaking in terms of transitioning into MR since most AR devices currently integrate additional sensors or external cameras that provide information about the surroundings, leading to potential uses on this area.
- AR is best suited for tasks that are applicable directly to the real world, such as a manufacturing factory since you can either interact with virtual appliances or with real ones by being able to see through both realities.
- AR devices are usually better for using them in a constant and mobile basis since they do not hide reality as VR does.

Therefore, we can assume that AR and MR will play a very important role as a present and future technology, even beyond VR. VR is great for generating completely virtual and different environments, non-real world related, and will continue to grow and develop but in some specific environments such as gaming. However, AR can expand to way more sectors and increase productivity in a more extensive way than VR could.

IV. Background

As described in the introduction section, AR is proven to be an unmatched platform to perform certain tasks or to enhance already existing ones. Augmented Reality's reach, however, is still somehow limited by the devices used to create this AR environment. Since these type of devices are created in a way that their intrusion to the user gets to minimum levels, so that they can immerse themselves in the features that they offer, they are often not powerful enough so as to actually generate a whole layer of virtual reality over the complete field of view of the user yet. The screening techniques used by most of these devices are still limited to a fragment of a human's actual sight, where you can feel the endpoints of that fictional reality.

In this project, however, we will make use of what are commonly known as Large Scale Immersive Holograms, whose objective is to simulate real-scale large holograms such as buildings, sculptures, natural elements and other sorts of large objects so that the user feels or perceives them as real. The use of AR for this purpose is highly recommendable as one of the potential uses of these type of holograms is to actually assess whether or not a potential product/building/feature would fit or be appropriate to be applied to the real world. For example, an architect may use LSIH to accurately check how a new skyscraper would fit in a given space, a manufacturing company could use LSIH to pre-emptively know if a certain layout of the factory's machines would work out in a given space, etc.

For this particular project we will specify the environments and devices that will be used. As the AR device we will be using Microsoft's HoloLens (v1), a standalone AR/MR device with the following features:

Software	<ul style="list-style-type: none"> • Windows 10 / Windows Mixed Reality
Optics / Display	<ul style="list-style-type: none"> • 2.3 megapixel widescreen see-through holographic lenses (waveguides) • 2x HD 16:9 light engines (screen aspect ratio) • Holographic Density: >2.5k radiant (light points per radian) • 1x 2.4-megapixel photographic video camera Automatic pupillary distance calibration
Sensors	<ul style="list-style-type: none"> • 1x IMU (Accelerometer, gyroscope, and magnetometer) • 4x environment sensors • 1x energy-efficient depth camera with a 120°x120° angle of view • Four-microphone array • 1x ambient light sensor
Processors	<ul style="list-style-type: none"> • Intel 32-bit (1GHz) with TPM 2.0 support • Custom-built Microsoft Holographic Processing Unit (HPU 1.0)
Memory	<ul style="list-style-type: none"> • 2GB RAM
Storage	<ul style="list-style-type: none"> • 64GB Flash
Wireless	<ul style="list-style-type: none"> • Wi-Fi 802.11ac wireless networking • Bluetooth 4.1 Low Energy (LE) wireless connectivity

Table 1: Microsoft HoloLens (v1) hardware specifications



Figure 2: Microsoft HoloLens (v1) device

To actually convey samples of LSIH to the HoloLens device we will be using Unity3D as our development/programming environment. Samples being tested will be based off a LIDAR scan of Las Olas Boulevard, Ft. Lauderdale, FL.



Figure 3: FARO's LIDAR scanner device

A LIDAR (Laser Imaging Detection and Ranging) scanner provides very accurate 3D scans of the areas being targeted and can export them into several formats so that they can be recreated through 3D modelling software. For the purpose of carrying these extensive scans to the actual Unity3D environment we will be using several tools, including:

- **FARO SCENE software.** This software is able to read a LIDAR scan file and generate a point-cloud model out of it. It also allows to generate those models with a variable degree of fidelity (i.e. number of polygons) and to select which parts of the scan should be used or not.
- **Meshlab.** This tool allows us to work with the mesh objects created by the SCENE software and to polish certain details before translating the object itself to Unity3D.

In order to perform some time-intensive tasks such as the rendering of those LIDAR scans, professor Hajek has built a rendering server at IIT's Rice Campus, remotely accessible, so that those tasks are performed by an external server working 24/7.

V. State of the Art

One of the major fields in which AR can prove to be very useful is the manufacturing and design field. This field has always been a huge asset to almost every nation's economic development, with the sector being more and more demanding these days. Over the last decade, digital manufacturing has become the go-to standard in this industry, as it allows manufacturing engineers to minimize the mistakes that could happen in the line of production to a very small extent, and in a fraction of the time needed before. However, AR technologies that combine both realities in a context-sensitive way can make room for techniques in which, in combination with a human being assisted by this technology, can help provide efficient and complementary tools to assist the manufacturing industry. This is heavily discussed by Nee et al. at [1], who dedicate their first part of the paper to explaining the availability of devices in the market matching the different AR sensorial experiences (e.g. Head-Mounted Displays or HMD for visual stimulations, gloves and other types of devices for haptics and tactile stimulations, etc.).



Figure 4: Example of AR assistance in a manufacturing process

Different software to support these devices in generating those renders or those stimulants that overlay with our real world is also discussed. The main focus of the article though is how deep the current research for AR (and to a lesser extent, VR) is, stating that nowadays most of the AR design efforts and existing research is based upon the prototype designing phase. This is clear for some big industries such as the automotive one, where AR could become a great asset to prototyping new pieces, car bodies, test angles, twists, even physics themselves, without the need to be interacting with something real which could be more time and money consuming. It also states how AR can have a big impact in a critical sector such as the medical one. AR has successfully been implanted in several different hardware devices, such as HMDs or even robotic arms to help assist surgeons with their task, being able for example to display critical measurements such as vital signs, ECG graphs, identify where a tumor is located, etc.

One among the other multiple tasks who could benefit a lot from AR is FLP, or Factory Layout Planning. Before the emergence of AR, there have been multiple VR approaches to take on this issue. Siemens, UGS and CAD Shröder, among others, have developed solutions in VR to plan and design factory layouts before its actual implementation. However, the main issue with this solution being developed in VR is that as it is simulated in a completely virtual environment, every miscalculation or deviation from the reality could significantly alter the final result, whereas if done in AR, one could just simply overlay objects, machines or whatever piece is necessary for the layout of the factory over an empty place to see if it fits, how could they be placed, etc.

Augmented Assembly is also discussed, in which within an Augmented Environment virtual objects can be created and manipulated to help assist with the assembly line. Within this

category, two main distinctions can be made: operating those systems without a special need for any kind of controlling gadget or instrument (namely gloves, mouses, trackpads or other tools) and those who do make use of them. The natural implementation though can be done by just implementing a natural human computer interaction interface, where human bare hands can be used as interaction tools with the AR overlay.

There is, however, a series of challenges and difficulties that AR technology has to face before being a completely reliable system for production. These can be divided into several categories, with the main ones being accuracy, registration, latency issues and AR interfacing issues. Accuracy makes reference to tracking spatial references correctly and scaling appropriately when overlaying information upon the real world. Registration is somewhat tied to the previous issue, and focuses on being able to place an object correctly in an augmented space. There was a solution presented in [2] which could potentially eliminate this errors giving a centimeter level precision both spatially and temporally. Latency issues makes reference to the obvious statement that things overlaying upon the real world need to be perceived as if they were on real time almost, especially if we consider the particular implication of a production line which must be both effective and fast. The last one, AR interfacing, refers to the issues that could potentially emerge as an AR interface should be intuitive, informative and immersive. This is a step-up from traditional user interfaces, in which these goals should also be pursued but the non achievement of any of them is not usually flagged as a critical issue.

As the main focus of this project is for it to be developed within Microsoft's HoloLens, other research papers and written resources, such as previous years' alumni work has been reviewed, which are written within this framework.

In [3], the actual architecture and design of the device itself is discussed, highlighting the fact that it is indeed the first mixed reality untethered (meaning not having to be connected physically to any other device) device to exist. It describes all of the optical subsystems that make the Mixed Reality happen, as well as all the electronic devices needed to support them: a custom GPU named HPU, a full logic board running Windows 10 as its OS, a broad Inter Pupillary Distance (IPD) range, IR sensors and LEDs, cameras, display engines and TOF depth map sensor. This paper concludes stating that Microsoft's HoloLens have a great degree of both comfortability (as they are untethered, have large IPD coverage, small size and lightweight, a good balance and pressure point and great brightness and contrast) and immersion (great Field of View range, world locked spatial audio, and an accurate gesture sensing).

Research found in [4] talks about how Microsoft's HoloLens can be used to produce large scale immersive holograms, which surpass the category of just overlaying simple objects over the real world, and instead are capable of modelling entire buildings, machinery or very complex and large figures. Different techniques to achieve those are explained, ranging from a basic approach in generating those bodies or figures in Matlab, to computational architectural forms using Rhinoceros 5 with the Grasshopper plugin through Unity for HoloLens. Although the research found here does not dig into interaction with the projected objects, it does state that the degree of immersion achieved by the HoloLens is quite good.

This large scale immersive holograms described by the paragraph above matches with references [6], [7], [8] and [9], which were research projects conducted on this topic at IIT from previous years students. A broad range of topics within the large scale immersive holograms theme are covered, such as human computer interaction, modelling a 3D model of an existing place, optimizing that model to render correctly and with a good performance within the HoloLens' hardware and interfacing.

State of the Art Conclusions

AR is an emerging technology that has already been proved very useful for its implementation in certain industrial sectors, such as the manufacturing and design one, and has the potential to be implemented in even more sectors. Existing solutions written and developed in VR have also been documented to improve, in certain cases, if applied to AR instead. A clear example of this is found in the Factory Layout Planning needed to setup almost any manufacturing or production building. VR covers, hides reality and transports the user to a completely virtual environment in which to test their desired layout. With AR, however, we could create a Mixed Reality environment in which the two realities would overlap, thus being able to precisely map each of the machinery the factory will need, in real time and without having to simulate the environment or building in which the factory will be set.

Although there are plenty of hardware devices that have the capability of delivering AR technology to the users, as the main device we will be using is Microsoft's HoloLens, and given that the main topic we will be covering is the use of Large Scale Immersive Holograms, we have

ensured that the device itself has the optimal technical specifications and hardware capabilities to deliver a exceptional performance for this task.

VI. Objectives

The objectives to achieve throughout this project are outlined in several subsections within this section. We have several stages of project development, namely:

- LIDAR scans' feasibility studio to act as LSIH
- Point-cloud rendering techniques with FARO SCENE
- Mesh objects polishing and adaptation to Unity3D with Meshlab
- Unity3D setup and performance tests with HoloLens
- Spatial Mapping concept proofs
- Future technologies: how can LSIH be deployed in the future

Overall, the main objective of this project is to focus on how LSIH can be created and integrated within an AR environment with the current technologies and limitations, how “immersive” and “realistic” the overall experience is, and how those two parameters are likely to increase with the application of newer technologies. Below we will describe what every of these subsection aims for.

Feasibility of LIDAR-scanned objects to act as LSIH

This section of the document will handle whether or not LIDAR-scanned objects are suitable to act as Large Scale Immersive Holograms. We will outline how a LIDAR scan looks like in its entirety, how it can be used to mesh the whole scan or just some elements and check the fidelity and realistic feel of its contents. This will also address how suitable the LIDAR scanning technology is to potentially build LSIH based in existing real-world content.

Point-cloud rendering technique analysis with FARO SCENE

This section of the methodology will address what tools does FARO SCENE offer when it comes to generating the point-cloud graphic necessary to later on create a meshed object directly from the LIDAR scan files. Which different types of rendering techniques do exist, which of them are available within the software and a comprehensive analysis on what each one of them offer. This analysis will dictate which of the techniques or specific parameters for those techniques is most suited to approach this particular case.

Mesh objects polishing and adaptation to Unity3D with Meshlab

Once a mesh object is created from the subsection above by using FARO SCENE, the objective here is to polish the details and remove unnecessary meshes or “noise” generated by cropping the mesh object directly from the point-cloud render. Also trying to remove as much workload as possible for the HoloLens GPU to handle, i.e., trying to reduce the polygon count without making it too obvious on the visual aspect, smoothing surfaces, etc. After the polishing and smoothing process is done, export the mesh object as a file that Unity3D can easily import such as an *.obj file.

Unity3D setup and performance tests with HoloLens Device/Emulator

One of the goals of this project is also to understand and learn how a Unity3D project should be configured in order for the HoloLens applications to be deployable and work. This will cover essentially all the dependencies and libraries that an Unity3D project needs in order to work within the HoloLens device as well as specific parameters related to the device itself such as the depth of view, field of view, etc.

In this phase of the project the main goal will be to assess how good several deployments of different qualities look once deployed to the device, and to do it from both the subjective (i.e. testing out the device and judging if it does look good or not just by perceiving the holograms) and the objective perspective (testing some objective parameters such as delay, frame rate, etc.). This assessment will try to search for a balance between a realistic experience and a smooth experience performance-wise, accounting for all the hardware limitations we might face.

Spatial Mapping concepts and tests




Figure 5: Spatial Mapping technology example

Spatial Mapping is a technique used by AR technologies (actually closer to MR than to AR) to use real-world data by scanning its surroundings, for example, with infrared sensors, to build a mesh of how the objects and elements nearby look like from a height, shape, dimension and volume perspective. Depending on how intensively this scan is performed, the resulting mesh will be more or less detailed as seen in Figure 5.

This section of the methodology pursues the goal of analyzing and testing how Spatial Mapping would help from a LSIH-deployment perspective to make the experience feel more realistic. Knowing how a particular hologram could be placed according to what the user is currently seeing in the real world can make the experience much more credible if done correctly. We will therefore deploy and test how some basic figures interact, scale and adapt to its surroundings and assess whether or not this would be suitable for LSIH.

Future technologies/devices for LSIH

Lastly, the objective for this project is for it to be scalable and a reference looking forward for new technologies and devices that can enhance the user experience. We will then discuss what the future looks like for AR, MR and LSIH and how new technologies will help make the overall perception better.

This section will primarily focus on the discussion and feasibility of the following technologies and devices:

- Microsoft HoloLens v2 (scheduled for late 2019)
- External GPU processing
- Cloud technologies: cloud-rendering, streaming services...

Each of these subsections will describe what these technologies or devices bring to the table so as to improve the experience and in which way they do so.

VII. Methodology

In this section of the document we will describe, for each of the subsections below, what the methodology procedures have been to achieve each one of the objectives described in the previous section. Here we will include a detailed insight on what each of the tools involved in this project are capable of, how have they contributed to the project and documentation so as to facilitate the continuity of projects in this specific area.

Most, if not all, of the software-related tasks are performed, as mentioned before, in a remote server located at IIT's Rice Campus in order to keep resource-intensive tasks as much time as necessary without having to dedicate a personal computer to those tasks.

Feasibility studio and analysis of LIDAR scanned files to act as LSIH

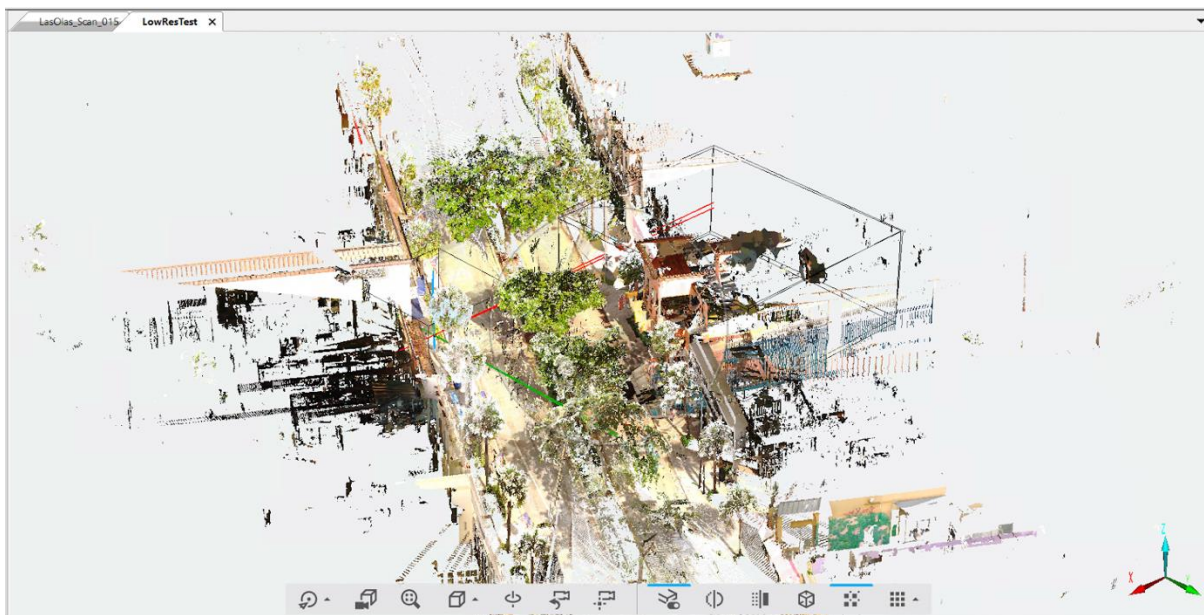


Figure 6: One partial LIDAR scan as seen on the FARO SCENE software

This section of the methodology is going to focus on how viable LIDAR scans would be if used to generate mesh objects to be displayed as holograms (LSIH particularly) from a visibility and perception perspective.

These scan files can be opened by using the FARO SCENE software, provided by the same company from whose LIDAR devices the scans come from. As mentioned before, these scans belong to Las Olas Boulevard, Ft. Lauderdale, FL.

Once the scan files have been opened by the software we are shown a preview of actual footage caught by cameras instead of LIDAR-based scans so we get a feel of what the environment looked like from a photography perspective. This would actually be quite similar to what Google Maps does with its Google Street View function. After that, the software must process the scan files before displaying the actual scans in a 3D fashion within the application itself. Once these processing tasks are done, we can open a view combining or just selecting one of the several scans performed, and we would obtain a view similar to that shown in Figure 6.

As we can perceive, the first impression is that it looks a little bit messy and noisy, and that it does not resemble a real street nor does it feel like an accurate model. However, we must take into account that this is not the final result, this representation only shows what the actual scan data has captured in terms of points, shapes, sizes, textures and colors. There are noticeable empty gaps between some of these objects or even within them, but once we build a mesh object from this data those gaps will be filled and smoothed out. So even though the first impression might not be a good one and therefore we could think about this models not working out to fit as LSIH,

through postprocessing and adapting these scans to resemble a higher level of fidelity we can still achieve great results.




Figure 7: Comparison between what we see in the LIDAR scanned 3D view (top figure) and what Google Street View actually provides (bottom figure)

On Figure 7 we can see that the comparison between the view that the scan files provide and the actual view from Google Street View are not that further apart. In fact, the LIDAR model is still very accurate if we do not mind the gaps between certain textures that, as said before, can be filled and smoothed out. So besides the inner building details not being too accurate (this is mainly due to how a LIDAR scan works, as shown in Figure 8) as LIDAR scanning light pulses

would reflect on the building's facade, and therefore the interior of said buildings must be scanned separately if we wanted the LSIHs to be immersive enough so as to let the user enter the virtual buildings.




Figure 8: Working principle for LIDAR-based scanners

Therefore, for the moment being we can conclude that these scans are not by any means realistic enough by themselves, but can be treated in a way that makes them feel realistic and thus could be suited to become LSIH.

Point-cloud rendering techniques available in FARO SCENE

In this methodology section we will address the step-by-step process of going from a raw LIDAR scan file to exporting a partial fragment of those scans as a mesh 3D object that we can work with, as well as the techniques used by FARO SCENE to achieve those results.

Importing the scans

First of all, we must import all LIDAR scans to the FARO SCENE software application. LIDAR scans generated by FARO equipment usually have a *.FLS extension. After locating and selecting those files, we should have them successfully imported into our system as shown in Figure 9.

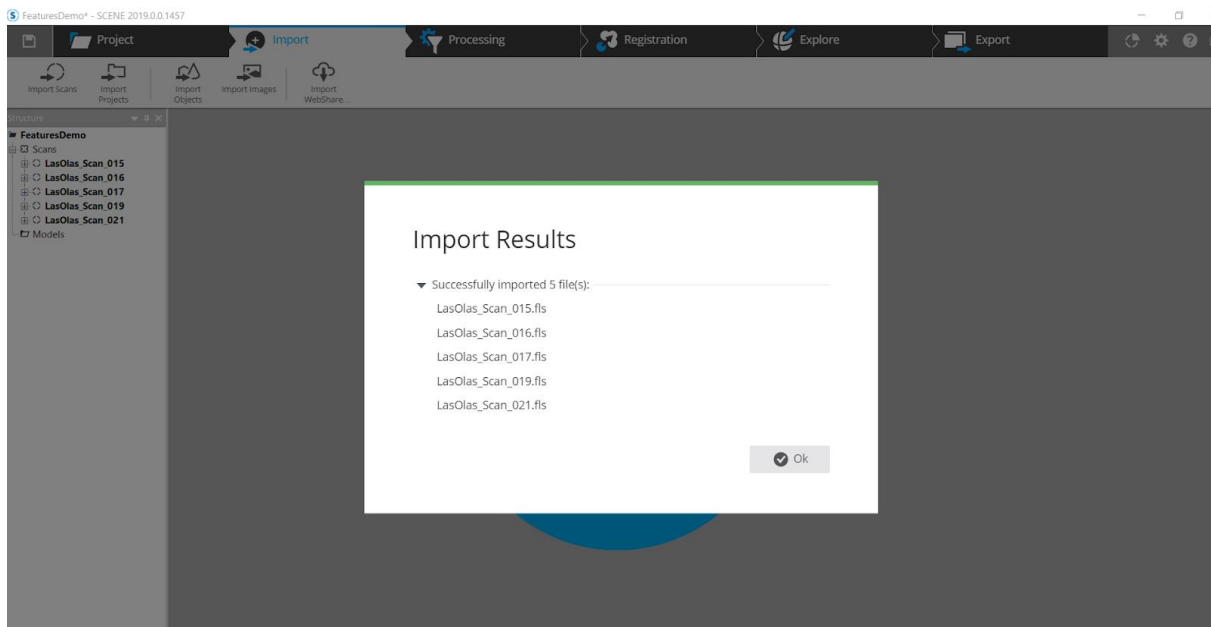


Figure 9: Scans imported into FARO SCENE

After those scans are imported and before we can proceed to visualize their 3D equivalent, we must process the scan files to generate the 3D model.

Processing the scans

In order to correctly process the scans we must first select which ones we will be processing and set/configure certain parameters to control the outcome of the processing task. Some of these parameters include:

- **Colorization:** By default this will generate just the 3D model of what we have in the scan files, but there is the option to colorize the model based on the pictures that the LIDAR scan takes in conjunction with the scan itself.
- **Filters:** We have several pre-processing filters available to generate a more accurate model, those being:
 - Dark Scan Point Filter: This removes points obtained from the scan that do not meet a minimum threshold of reflected light. This could potentially remove noise from the model itself if we stop considering reflections that are not relevant for the actual shape we are trying to scan.
 - Distance Filter: This filter can basically crop the generated 3D model to a fixed distance from start to end, and remove all points exceeding that range.
 - Edge Artifact Filter: This filter smoothes out or straight out removes edges on shapes found on the scan files in an intelligent way, preserving surfaces such as floors, ceilings and walls so that the overall model looks more polished.
- **Find targets:** This option is mainly used for the processing to be focused in a particular selection done over the scan files. These selections should be done by hand but this option can focus on looking for manually introduced markers, spheres or planes to act as selections.

Registering the scans and creating mesh objects from the 3D view

Before we proceed to the next step we can register the scans within the same scan cluster to have them grouped. This can be done manually by indicating how every single scan overlaps with the others or in an automatic way by letting the software decide how to handle them.

Now we can switch to the “Explore” tab within FARO SCENE and proceed to open the 3D model view of the scans. We have a pretty intuitive move-around tool where we can fly over our model or rotate our view to reach every angle we want to check.

Our main purpose within this application is to convert this 3D scanned view to actual objects to be displayed as LSIH. There are several approaches to achieve this goal, including directly rendering the whole 3D scan as a single, big mesh object, but that comes with a lot of downsides considering that the 3D scan is not clear and clean enough to do so. There are things within the scan that we can consider trimming down, especially knowing beforehand the hardware limitations that we will face for these holograms to be rendered. Therefore, the approach chosen in this project is to start with small scale single buildings and react to the performance of the device itself: if it is considered good, we can add more details or elements; if it's not, then we look for some other ways of improving them.

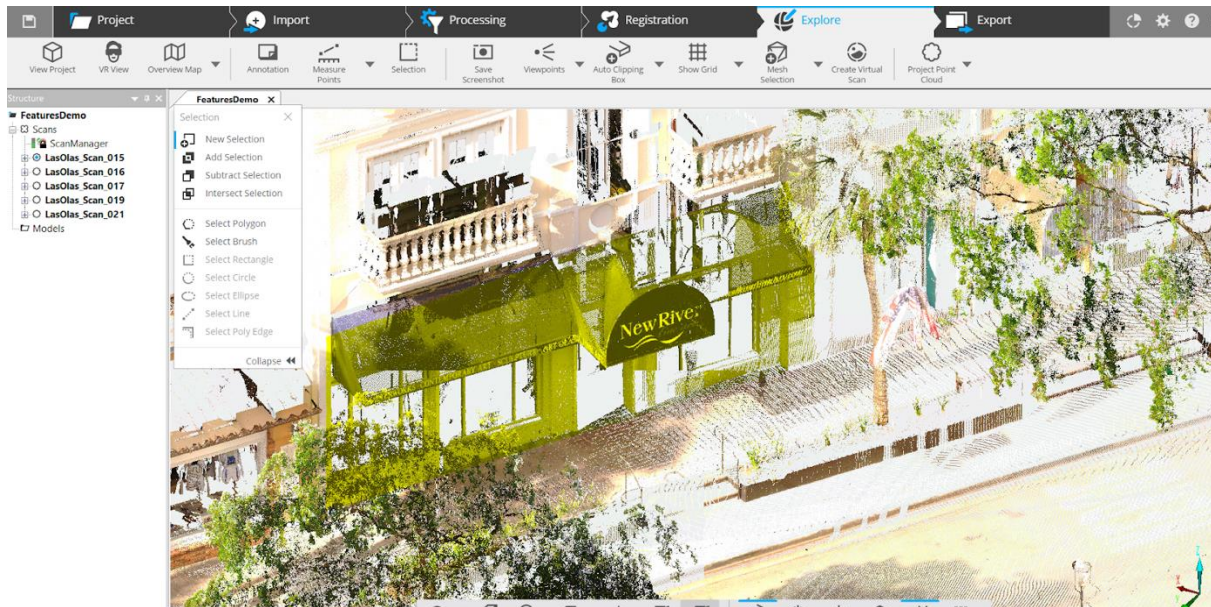


Figure 10: Selection tool used to select one of the buildings of the 3D scan

In Figure 10 we can see how the selection tool can be used to highlight certain parts of the 3D view to later convert them into a mesh object. The selection tool offers a broad range of modes to make the selection: polygons, different shapes, a manual brush... Once selected, we can add more to our selection, remove areas that we selected erroneously, intersect them if we are taking the selection from multiple angles...

Once we finish our selection, we can go to the “Mesh Selection” option within the top-bar menu, and we will find the prompt shown in Figure 11.

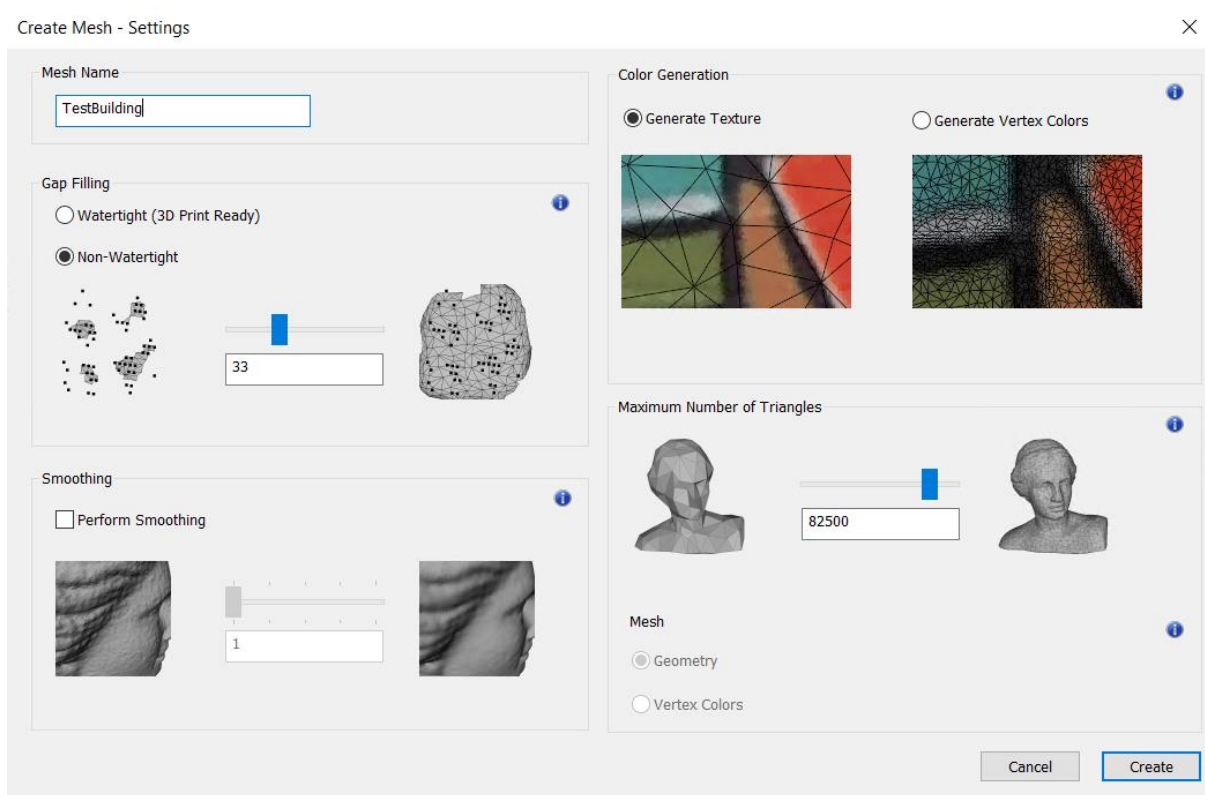


Figure 11: Mesh selection options available from FARO SCENE software

Below we will proceed to describe every single one of these options.

- **Gap Filling:** This is the technique used to fill the gaps between the point-cloud generated by reading the LIDAR scan files. This is what will return us an actual consistent object instead of a messy building skeleton. We have two options within this field.
- **Watertight:** this means that the mesh created will be completely enclosed to avoid gaps if the mesh object we want to create is or can potentially be sensitive to water. It also means that having no gaps it is almost guaranteed that it would be 3D-printable, as some 3D printers build the object layer by layer and can not cover gaps depending on their position.

- **Non-Watertight:** this means that not every single gap between points has to be necessarily covered. In fact we are presented with a very descriptive slider that lets us adjust the threshold on how close should the points be in order for the algorithm to join them together.
- **Smoothing:** This technique basically reduces the harshness or pointy ends that the resulting model could have. It tries to resolve the mesh to a smoother surface that can cause an object to lose detail or become smoother, so this is an option to be used with caution.
- **Color generation:** This references how the coloring will be added to the mesh object itself. Colors could be generated through textures overlapping the existing mesh polygons or be generated through Vertex colors. This last technique assigns a color for each vertice and calculates the color of each polygon pixel based off the colors their vertices have.
- **Maximum number of triangles:** This basically determines the geometrical complexity of the mesh to be generated. The more triangles or polygons we allow our mesh to have, the more close to reality the resulting 3D object will be. However, this increase in quality is directly related to an increase in resource consumption when deployed in a device, as well as in the model's storage size, so a balance must be kept. As this can be easily post-processed, it is recommended to set for the maximum amount of triangles and then reducing them if necessary.

Mesh objects polishing and adaptation to Unity3D using MeshLab

Once our mesh object is exported (preferably in a *.ply format or similar so that the color information is kept) we can proceed to review it on MeshLab.




Figure 12: Unpolished, sample mesh object exported from SCENE to Meshlab

Through MeshLab we can perform plenty of actions with our 3D model including removing some aspects that were selected with the FARO SCENE selection tool but we do not want them to be into the mesh itself such as those tree fragments. By using the selection tools and the remove vertex option we can get rid of them.

After a little bit of polishing and just by removing unwanted elements that would only add more processing load we end up with something like Figure 13.




Figure 13: Building facade with unnecessary elements trimmed down

We have to keep in mind that as the interior of the building is not scanned properly due to how LIDAR technology works, we've decided to remove it altogether as it did not provide any accurate element of how the interior actually looks like and would only add more resource consumption. We could later on scan the interior or render our own version of the interior and add it afterwards behind the actual facade.

There are, however, some necessary adjustments that we must make before transitioning this model into a potential LSIH. The LIDAR scan worked pretty well but still missed some points in the right-hand side of the awning, most likely due to the LIDAR scanning device being placed on the left and being unable to reach that particular spot. This means that there is a part where no awning is appreciated, and instead a huge gap is shown. Even though SCENE offers a gap-filling feature this is not intended for this kind of usage as if we were to try to recreate this building as realistically as possible, we would need to take a LIDAR scan again from the right-hand side of the building to overlap both scans and have the facade from both angles. In fact, we can see that raising the gap-filling feature by a significant margin can lead to unwanted elements being enhanced or introduced:




Figure 14: Differences between gap-filling settings

We can appreciate that there are some noticeable differences between the two models, mainly how the showcasing window frames are displayed. With the low gap-filling threshold we get these frames to be more reliable and realistic, and with a higher gap-filling threshold we find them to be unnaturally enlarged. The same happens with the overall facade, as it looks more rounded than it should be in reality.

With this said, gap-filling is a very interesting parameter to look out when extracting a mesh out of the point-cloud generated model and there isn't always an ideal threshold to get the model to represent reality in an accurate way, so trial and error method is mostly preferred to reach the ultimate design.

Now we will discuss the post-processing options that Meshlab provides for these objects and how these options may prove useful for when we deploy the object into the HoloLens device.

Smoothing and remeshing techniques

Meshlab offers plenty of filters to smooth out an existing mesh as well as repairing or re-meshing the object. We will discuss some of them:

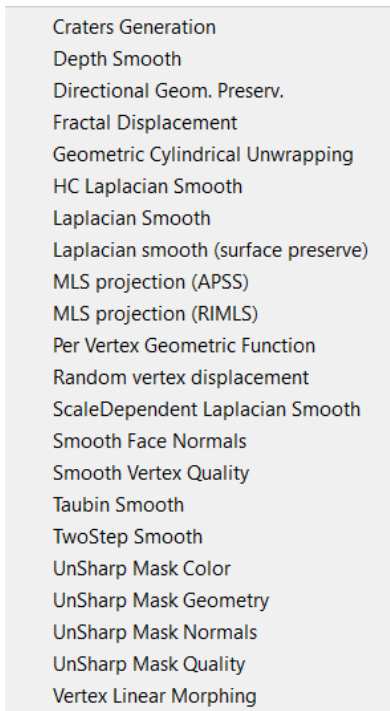


Figure 15: Some of the smoothing filters that Meshlab offers

- **Laplacian Smooth:** This filter is intended to reduce the overall noise a mesh can have, or even enhance or exaggerate some of its features by applying a negative coefficient to the filter parameter. It tries to preserve the original geometry while taking rid of this noise and can be done iteratively, which means that the mesh would pass the same filter a number N of times.
- **Depth Smooth:** It applies the same principle used in a Laplacian smooth filter, but it performs it just from the designed point of view, meaning that not the entirety of the mesh has to be affected by this filter, just the part designed to be viewed by the user.

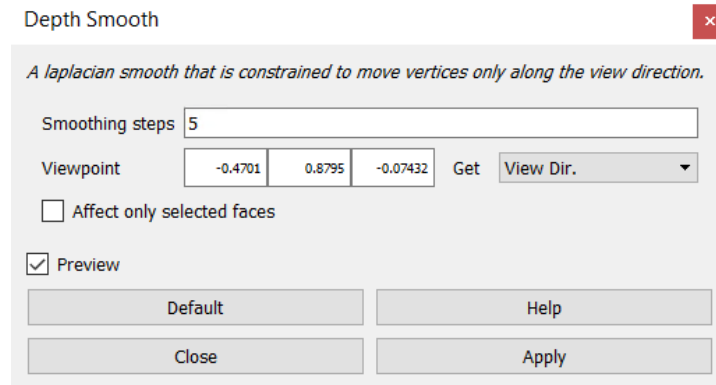


Figure 16: Depth Smooth parameter setting window: we set the viewpoint and the iterations for the laplacian smoothing

- **Smooth Face Normals:** It performs a smoothing process without changing the vertex points set for the mesh. It just tries to smooth out the resulting faces across all vertex points forming the mesh object. Results are milder but does not alter the mesh structure.

For the remeshing techniques more focused on toning up/down the quality and therefore the performance of the 3D object we will mainly make use of one technique called Quadric Edge Collapse Decimation that is able to reduce the number of polygons a mesh object has.

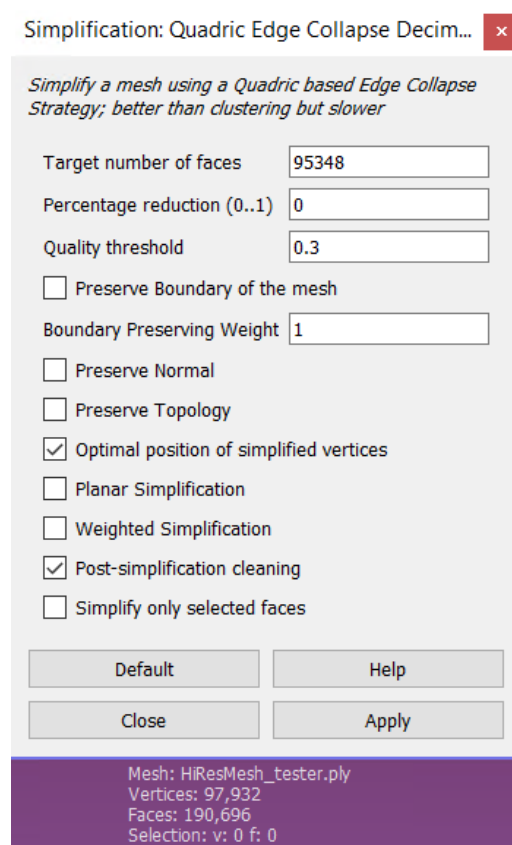


Figure 17: Quadric Edge Collapse Decimation filter configuration window




Figure 18: Mesh before and after a polygon reduction

We can notice here that even after a significant polygon reduction the overall structure of the facade is still the same. This is because in this particular mesh object most of the vertex points are not used to define the actual structure of the facade itself but to mark where the colors should be: there are flat surfaces that normally would require just a small amount of faces/polygons that

actually have lots of them just to accommodate different colors and color shapes within the same surface. This is especially noticeable in the logo displayed on the awning, where the text is more or less legible within the original mesh and then becomes blurry after the polygon reduction.

Unity3D setup and performance tests with HoloLens Device/Emulator

Once we have the objects we want to deploy to our HoloLens device, we need to setup an Unity project suited to be deployed to the device. This section of the methodology will cover how to deploy a project to Microsoft's HoloLens from a fresh blank Unity3D project as well as the requirements needed for that purpose. Later on, this section will discuss how are we going to test the performance the device outputs when rendering our LSIH.

Requirements to deploy an Unity3D project to HoloLens

Before proceeding to deploy a project to the device, we need to make sure that we are running the following list of requirements:

- Windows 10 as our main OS, to the latest update
- Visual Studio 2017 Community Edition or higher.
 - Community Edition's license is free as long as you have a (free) Microsoft Account.
 - Within the Visual Studio environment, we need to install the following additional packages:
 - Desktop development with C++
 - Universal Windows Platform (UWP) development
- Windows 10 SDK, latest revision

- (Optional) HoloLens (1st gen) Emulator.
 - To use the emulator, it is a requirement that the system supports Hyper-V
- Unity3D LTS latest version (currently 2018.4)
- MixedRealityToolkit assets and packages

Configuring the project settings

Once the above requirements are met, we will need to configure the settings of our newly created Unity3D project. Most of these settings are accessed through the Edit > Project Settings window.

First of all, we would want to change the quality of the project for UWP targets to be as low as possible. This is due to the hardware limitations that the device has.

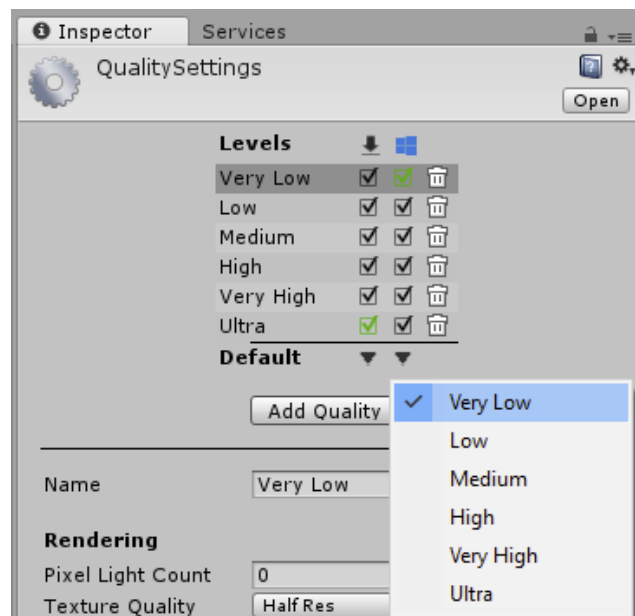


Figure 19: Setting very low quality for UWP

Furthermore, we need to let Unity know that we are deploying the project to an AR/MR device and that therefore, the application to be deployed should try to create an immersive view instead of a 2D plane. This is done by enabling Virtual Reality Support on the project by targeting Windows SDK. This is under the ‘player’ tab in the same window as before:

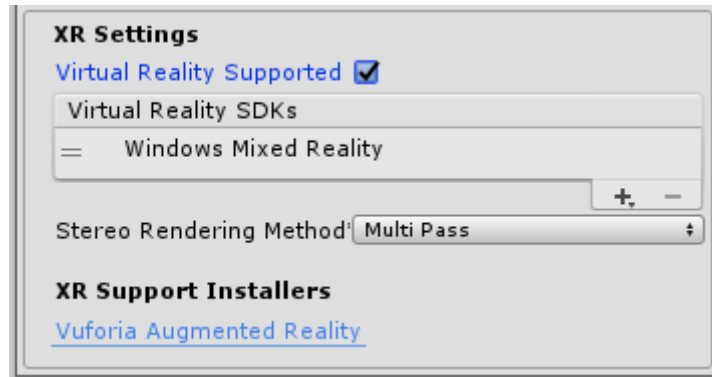


Figure 20: Allowing Virtual Reality Support in our project and targeting Windows 10 SDK

Next, we need to ensure that the .NET configuration is set up correctly. This configuration option is in the same pane as the Virtual Reality Support one, under the “Configuration” section. We need to make sure that the scripting backend is using .NET:




Figure 21: Setting the Scripting Backend to .NET

With these settings we should be all set to deploy the project once finished.

Deploying the project to Visual Studio

Once our project is finished and ready to be deployed, we need to follow the following steps in order for it to be compiled to be opened in Visual Studio.

First we should open the **File > Build Settings** window. This will prompt us a window where we can configure the deployment settings for the project. The first thing we should do here is to add every scene that we want to render. If we want to add every scene we have opened in the editor and that we have worked with, there is an 'Add Open Scenes' button which will automatically include them.

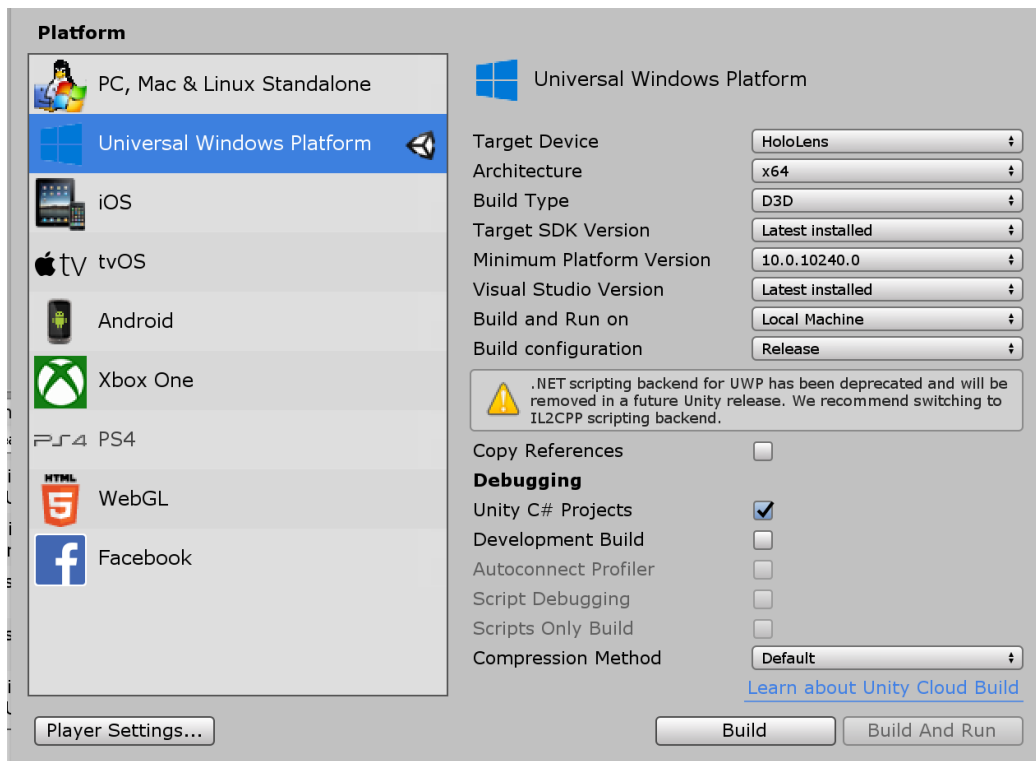


Figure 22: Build window settings

After that, we should make sure that we target the Universal Windows Platform as our deployment platform. This is done by selecting it on the left pane and then clicking the ‘Switch Platform’ button if we were not already using it. The Unity logo to the right indicates which platform we are using at that moment.

Then we want to configure the rest of options as follows:

- Target Device would be HoloLens
- Build Type must be D3D
- Target SDK Version should be the latest installed version
- Build and Run on should be set to Local Machine
- The ‘Unity C# projects’ under the debugging section should be ticked

Then we can click the ‘Build’ button and the project will start to compile. This project will not run by itself, instead, it will generate a Visual Studio project that we must open to deploy it to either the device or the emulator.

Now on the folder we selected for the project to be deployed on we should find a *.sln file that will automatically open with Visual Studio. Once opened and all files have been loaded successfully, the main toolbar will let us deploy it to our preferred device.

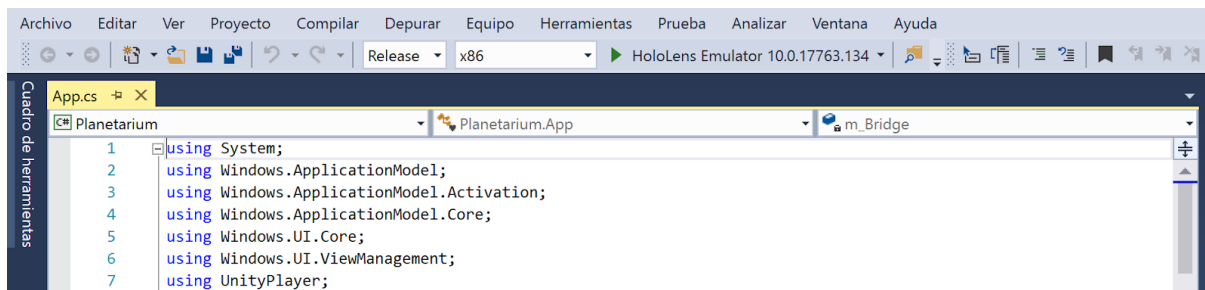


Figure 23: How Visual Studio looks like when deploying a project to HoloLens

We have several options here for the deployment:

- Deploy the project to the HoloLens Emulator, as noted in Figure 23
- Deploy the project to the HoloLens Device:
 - By connecting it via USB
 - By providing the IP address of the device. Must be on the same LAN.

Once all deployment settings have been correctly inputted, we can click on the play button so that the deployment begins. Visual Studio will compile the whole project and send it to the device or emulator.




Figure 24: Simple Project deployed on HoloLens Emulator

Spatial Mapping concepts and tests

Spatial Mapping is a technique that really transitions the AR features of Microsoft's HoloLens to a MR experience. It uses real-world data to create a more interactive experience for the user. The most common way these kind of devices get to interact with their surroundings is through sensors (which could be infrared, straight up cameras, laser tracking devices and in essence, any sensor capable of scanning 3D environments in a close-range setup).




Figure 25: How Spatial Mapping obtains its data to be applied later on

Within this context, a Spatial Mapping-capable device maps its surroundings by scanning them through its sensors (the HoloLens have 4 environmental sensors to scan the whole field of view of the user) and creating a virtual mesh to act as reference for where the objects are placed.

Later on, this data can be used to cover a wide range of purposes. By knowing how the real world around the user wearing the device is shaped, we could choose to place objects in a specific location or creating more realistic physics. We could for example create a spherical hologram to

act as a ball and to have it bounce taking into account how the real objects are distributed in space, so if the ball were to bounce over a table, it would do so earlier than it would if it were to bounce on the floor.

Essentially, this marks the main differentiation between AR and MR. In AR you can see both the real and the virtual worlds, but there is no interaction between them. You can use AR to overlay things over your normal view, to visualize UIs, LSIH or other types of static holograms, but there is no interaction whatsoever with the environment. In MR, however, everything can be suited and adapted to the real world. You know the data about your surroundings and how to apply that is just left to imagination.

So the next question would be how can this be applied to this environment. Since we've already predicted that the LSIH rendered in this project are not likely to have a good deployment in the HoloLens device due to their high demands in resource consumption, number of polygons and such, we will not be trying to apply Spatial Mapping to these specific 3D models but instead doing some field tests of simpler holograms to see how this technology could be applied in the future when processing power becomes less of a limitation.

We could define 5 main steps when trying to deploy and apply Spatial Mapping within a HoloLens project in Unity:

1. Scanning the surroundings
2. Visualizing the resulting scan as a mesh

3. Processing spatial mapping data to be used in the project: detect walls, ceilings, floors...
4. Working with placing actual objects/holograms within the scanned mesh
5. Adding extra details like occlusion, shaders...

We will focus the results section on this topic for this document in how we can scan a room with the HoloLens device, to especially focus in detecting its walls and ceilings, and having them act as a reference to place our LSIH objects.

Description of future technologies for LSIH

In order to better understand what types of technologies or improvements are expected to happen in the future of AR/MR and their LSIH application we need to know first the limitations that these technologies have nowadays.

The first and most obvious limitation nowadays would be the processing power that an AR/MR device has as most of them are supposed to be a lightweight HMD device and thus hardware limitations come naturally by the space restrictions that such devices have. There are technologies and hardware so as to make the rendering of LSIH seem like a breeze, but most of them are heavily dependant on space and power-supply and therefore incompatible with these kind of devices.

However, we can bring some solutions to this main issue by trying to externalize the computation process to a device other than the HMD itself. This is where we come with two main solutions:

- Using an external GPU setting (wired)
- Using cloud infrastructure to build our own external rendering solution

The second most common limitation would be the field of view that the actual devices have. This does have more to do with optics and laser display technologies than with computational power itself. There are significant improvements planned for the future of these devices, as Microsoft is going to boost their new iteration of their HoloLens with a broader and wider field of view which is supposed to cover twice as much area as their current iteration does. Still, that would probably be short when taking into account how much field of view a human eye normally has.

Therefore we can assume that in the near future, improvements will be made in these two specific areas to deliver a better experience with LSIH.

VIII. Results

This section will outline the results we have obtained for the following sections:

- LSIH deployment in HoloLens tests
- Spatial Mapping in HoloLens tests
- Future technologies applicable to LSIH and HoloLens

LSIH performance test results

The objective we set up at the objectives section of this document was to test whether or not these LIDAR-scanned LSIH models would be viable to use in the Microsoft HoloLens device with an overall good experience. Therefore the expectation was to check how many buildings or structures the device could handle at a time with a given quality and scale enough to appear realistic.

Before performing the tests by themselves, we came to realize that precisely because LIDAR scanned models are pretty detailed and in a good scale, and given the shortcomings of the HoloLens device in terms of hardware, that we would probably struggle to render just one building.

That is mainly why the building shown in the methodology section has been trimmed down to leave only the facade. The other reason is that the space within that building is currently empty because of how LIDAR scans work, as we would have needed to scan the building from the inside to get an accurate model of the interior. Anyway, the intention was to test three iterations of the same facade so as to have a fair comparison, with each iteration having a lower quality than the previous one. We wanted to test out these models:

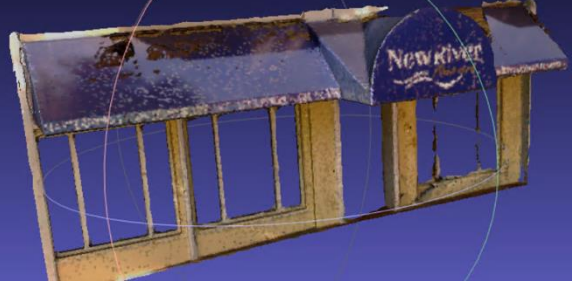

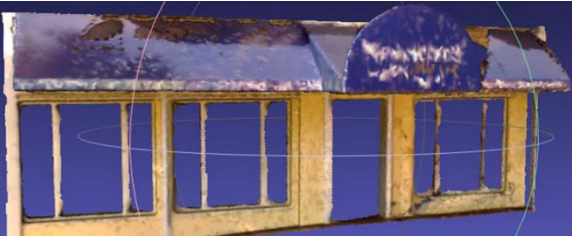
Model	Characteristics
	<ul style="list-style-type: none"> • Higher quality facade model • 130k+ vertex • 260k+ faces (triangles) • Color as vertex colors
	<ul style="list-style-type: none"> • Medium quality facade model • 90k+ vertex • 190k+ faces (triangles) • Color as vertex colors
	<ul style="list-style-type: none"> • Lower quality facade model • 35k+ vertex • 60k+ faces (triangles) • Color as vertex colors <ul style="list-style-type: none"> ○ Extracted texture as well

Table 2: LSIH models to be tested and their characteristics

It is appreciated especially in the awning's text how the quality and the polygon count drop appears more noticeable the lower the quality is. For the last model we also performed a texturization implying that in order to lower further down the polygon count, we let go of the vertex colors and substitute them for a simpler mesh and a texturized object.




Figure 26: The texture-based mesh imported into a Unity3D-HoloLens project

The results, however, were a little bit disappointing. Although every single of these objects compiled and were successfully loaded into the HoloLens device, the performance for the High and Medium quality ones did not result in a great experience. Instead, renderings were slow and laggy, and you could tell that the device was being tested to its limit. Even the HoloLens Device Portal wouldn't respond at some times and so we could not extract exact performance meters out of those two tests.

With the last one, however, we obtained an actually usable hologram, taking into account that the scale was 1:1 (so the building looked as tall as it would in real life) and disregarding a little bit of existing lag on the placement when moving around it, the results were not bad at all.



Figure 27: Rendered low-quality LSIH at IIT Tower 1st floor

As we can see the building is correctly deployed even though the quality of it could be better. It did not feel laggy at all and actually ran at the maximum FPS limit.

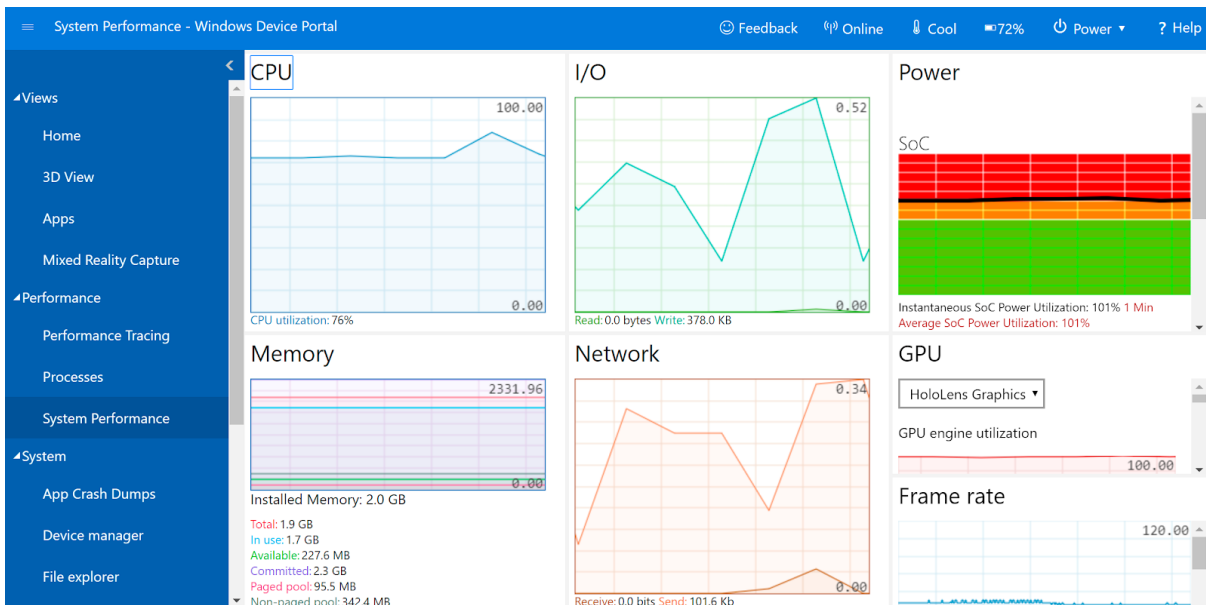


Figure 28: Performance for the low-quality texture based LSIH model on the HoloLens device

As we can appreciate in the performance parameters extracted from the HoloLens Device Portal, we are using all the instantaneous SoC power that the unit can deliver. This model is already consuming almost 1.7GB out of the 1.9GB of available RAM memory to the unit, using the Intel Atom X5 CPU to 76% of its capacity (when being static around the hologram) and the GPU (which is currently the custom-built HPU that Microsoft has embedded within the SoC) is at a 100% of its utilization. Even then, the framerate remains stable at 120FPS which is a significantly good number.

Spatial Mapping test results

As we described in the methodology section when trying to apply Spatial Mapping technology to a project we have 5 main phases which could in turn be simplified to two essential ones: scanning the surroundings and how that data is treated and used.

For the first part, the scan is done by the HoloLens device itself since it packs 4 sensors capable of tracking distance and volume to nearby objects, walls, ceilings and floors. Within the HoloLens Device Portal we find that we can track what the HoloLens is sensing in real time and even build a mesh model of our surroundings by just walking around.

The results are pretty accurate if we mind that this is quickly done by a HMD device that is not as powerful as other scanning technologies.




Figure 29: Top view of the mesh generated by HoloLens, at TS 2030 (SmartLab)



Figure 30: An inner view of the classroom. Objects such as chairs and tables are detailed as well.

We can see that for being a quick approach the results are quite good. This mesh is exportable as an *.OBJ format and therefore importable to the Unity3D environment. This is quite

good as we could use this mesh as reference even if we did not want to use Spatial Mapping at all. Meaning that by just performing a previous Spatial analysis on our working room we can have a scale within our Unity environment telling us how every hologram would be size-wise when compared to the rest of relatable real world objects and we could even work on object placement.

Now for the Spatial Mapping technology itself, Microsoft provides us with an already prefabricated script named `SpatialMapping.cs` capable of tracking the environment in real time. Within the options that this script has, we can enable turning on or off a guideline to help us visualize that real-time mesh tracking. In fact, we can see how the HoloLens would generate this mesh without the need to deploy any project at all. If we gaze somewhere with the HoloLens turned on, and we do the ‘air tap’ gesture, we’ll notice how the HoloLens generates a mesh scan of where we pointed our gaze, as noted in Figure 31.



Figure 31: Microsoft HoloLens performing a mesh scanning of where we pointed our gaze



Figure 32: Microsoft HoloLens displaying a real-time mesh layout when we deploy a project with the `SpatialMapping.cs` script with grid enabled

So as we can see by both pictures, the real-time mesh generation of the surroundings for the HoloLens device is pretty accurate and that data could be used to enhance our projects. One interesting thing about the Unity engine and this particular Spatial Mapping script that Microsoft provides is that we can configure colliders based on this mesh data. This means that we could enable several kinds of physics. We could make an object bounce based on this generated mesh, or we could deal with things such as positional-based occlusion.



Figure 33: Types of Occlusion available in Unity3D + HoloLens setup

As we can see by the figure above these different types of occlusion hide the object when it positionally collides with a real-world element. This creates a more realistic experience where virtual objects seem to interact with real ones. Depending on the type of occlusion applied, we can see the outline of the parts being occluded or not.

Deploying a test on the standard occlusion to our device resulted in partially good results. As the end results depend heavily on how well the spatial mapping mesh is rendered we will notice that sometimes occlusion does not happen exactly as we want it to.



Figure 34: Spatial Mapping occlusion example at the SmartLab. We can see the sun being partially occluded by the box over the chair

It is worth mentioning that even though the sun appears to have a poor occlusion there are other planets as well as the rock debris ring orbiting around it that are correctly hidden behind that box. So the results are noticeable and not bad at all.

Therefore, we can conclude that Spatial Mapping technology could contribute greatly once more processing power is available to these kind of devices as the transition from AR to MR greatly enhances the immersion and overall experience.

Future technologies and devices for LSIH

Microsoft HoloLens v2



Figure 35: Microsoft HoloLens v2 device

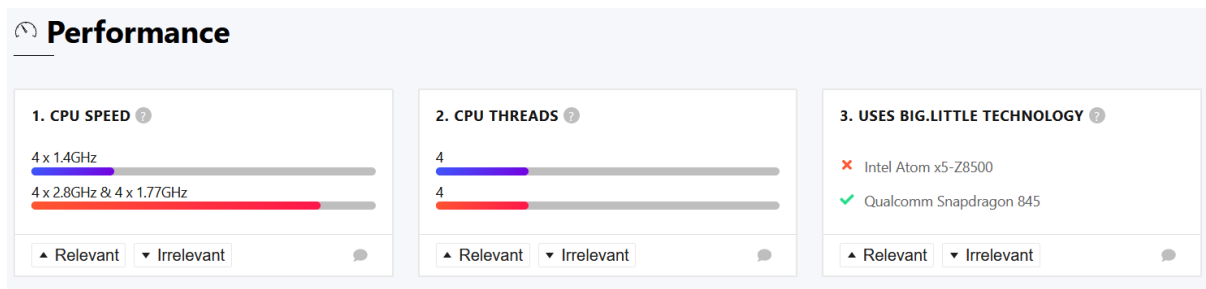
Microsoft announced this last February their new HoloLens v2 device which will be taking over the current version and add a significant number of new and improved features. Even though the full technical specs for this device have not yet been officially released by Microsoft, we have a couple of confirmed hardware features:

- **Resolution:** 2K 3:2 light engines in each eye
- **Holographic density:** >2.5K radiants (light points per radian)
- **Processor:** Qualcomm Snapdragon 850
- **Holographic unit (HPU):** 2nd-generation
- **Wireless:** 802.11ac (2x2), Bluetooth 5.0
- **Wired:** USB-C
- **Camera:** 8MP stills, 1080p video

- **Mics:** 5-channel
- **Speakers:** Built-in, spatial audio
- **Other features:** Eye tracking, head tracking, Windows Hello authentication, 6-degrees-of-freedom (6DoF) tracking

From these confirmed HW spec upgrades we can highlight two especially interesting features: the resolution now being 2K as well as the holographic density being upgraded, which will for sure improve how the holograms are perceived, and the processor now being a Qualcomm Snapdragon 850, which should include within its SoC (System-on-Chip) a Qualcomm Adreno 630 as its integrated GPU. This should result in a significant increase from the previous version Intel Atom X5.

We can find a full hands-on benchmark article referenced in [10] from which we can extract the following interesting data:



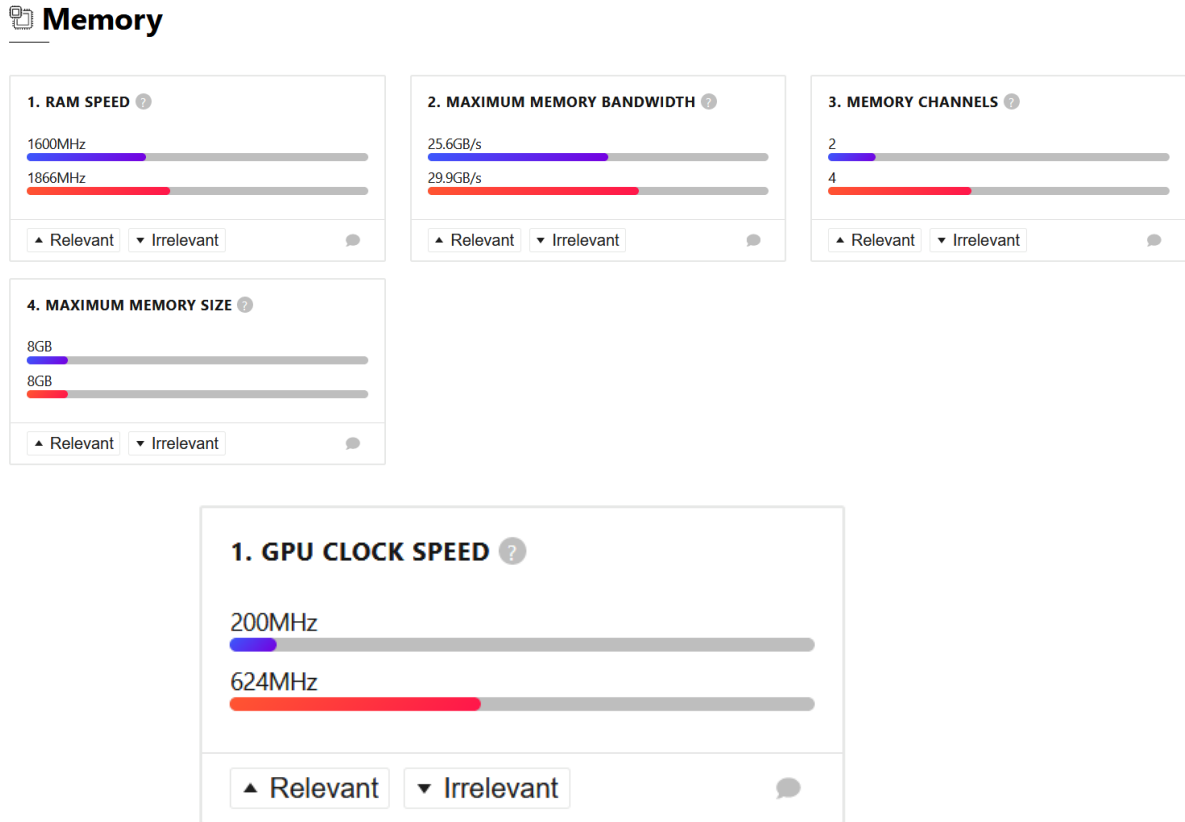


Figure 36: Benchmark comparison between Intel Atom X5 (HoloLens v1) and Qualcomm Snapdragon 845 (HoloLens v2)

This benchmark actually references the Snapdragon 845 instead of 850, but they are virtually the same SoC. The main difference is that the Snapdragon 850 comes already overclocked to add 0.1GHz more clock speed.

We can see that there are several aspects of the current HoloLens version hardware that were very limiting that are now being vastly improved. Clock speed for the CPU has been significantly increased by having 4 cores running at 2.8GHz (twice the current speed, even more if we take into account the 850's specific overclocking to 2.9GHz) and 4 co-processors running at 1.77GHz. RAM speed and latency has also increased to reach 1866MHz although its size remains

the same. We can see however how the integrated GPU from this SoC is, at least from a clock speed point of view, much faster and will allow for more complex tasks and holograms to be displayed.

From the official presentation at MWC 2019 in Barcelona, Spain, we can also recap a list of features and improvements done over the previous version. The most remarkable ones are:

- Iris recognition technology for authentication purposes
- The user field-of-view (FoV) has been greatly improved. While the previous generation had a 34-degree 16:9 aspect ratio field of view, the new HoloLens come with a 52-degree diagonal FoV which according to Microsoft sources is 2X the current area and most of its growth is now vertical. This is a huge improvement towards user experience which is, ultimately, what LSIH aim for. However, it could still be a little bit lacking when applying this data over a comparative layout as seen in Figure 37




Figure 37: Comparative for the FoVs of HoloLens 2, a VR device and the human vision

- Native hardware claims to be capable of rendering meshes or 3D objects of up to 100,000 polygon counts.
- Improved user gesture tracking system. While in the current device the front-facing environment-scanning cameras are used to take care of the user's hand gestures, the improved tracking system will allow for all of the user's fingers to be tracked, allowing interaction with more complex objects such as a virtual piano, or having more complex UIs set up. This tracking system would still require the user's hands to be within the tracking area of the device.
- This improvement comes with the addition of several UI new options such as buttons, sliders, etc.
- There is also a claim that this device is exclusively focused for businesses and business practices, disregarding other fields of application such as gaming. This does not mean that the HoloLens 2 won't be suited for those tasks but rather that its features will not be built with that in mind.
- The most important feature of all of them regarding this particular project is the announcement by Microsoft that they will allow some sort of cloud-rendering service through their Azure Cloud Computing Platform. This will be discussed in the next subsection of the results.

Cloud rendering/computing technologies

One of the main issues that we have faced throughout this project is the high resource requirement that Large Scale Immersive Holograms require. Even though the HoloLens 2 device has made significant hardware improvements over its previous iteration of the product, when we

talk about delivering a very immersive experience we must take into account that the objects or holograms we want to interact with must be very detailed and good-sized so as to accurately represent reality over a virtual environment. This means that there is almost no mobile processor capable of delivering the performance requirements these kind of holograms would require, at least with the current technologies.

Therefore, we have two possible options to try to overcome this issue. Both of them rely on forwarding the processing capabilities to another device. The first one would be relying in an external GPU unit to handle all the required graphic processing tasks. The second one is way more flexible and scalable towards the future and would consist on relying the computational power to cloud computing or cloud rendering services, such as the one that Microsoft has promised for this iteration of their HoloLens.

First option: external GPU unit

Since there is no official claims nor evidence that this kind of devices will be supported from Microsoft's end, we can theorize about their hypothetical contribution to LSIH deployment within HoloLens.

An external GPU unit is usually a case with a power source and an interface to let an external, independent graphic card (such as the ones used widely by the market in both gaming and graphic-processing intensive enterprises, as well as other recent uses such as mining bitcoins or cryptocurrency) act as the main source for computing graphics for the interfacing device.

This essentially means that all graphic processing is done in a different device and therefore the limitations are no longer those of the SoC/HPU that the HoloLens device packs. Instead, the limitations with this approach will be delimited by two main factors:

- The GPU being used by this external GPU device. There is usually a large margin here, as we could be using a low-tier desktop GPU (and still be better than a mobile processor's GPU) or a high-end enterprise GPU and the computational capabilities will be completely different. We can also find that we could do this external GPU device hold one or more GPUs, as both of the most popular manufacturers offer bridges that allow interconnection between two or more of their GPUs (SLI for nVIDIA and CrossFire for AMD).
- The physical interface and drivers used to connect the device with the external GPU. We do know that the HoloLens 2 come with USB-C, so we will take the assumption that this port is data-ready and that it works at the standard USB 3.1 Type-C data rate of 10Gbps. We could extend this to 40Gbps if we were to suppose that instead of just a USB-C connection we had a Thunderbolt 3 port, so we will outline both cases. This is very meaningful and probably the main source of bottlenecking as all the processed data should be sent to the HoloLens to be displayed.

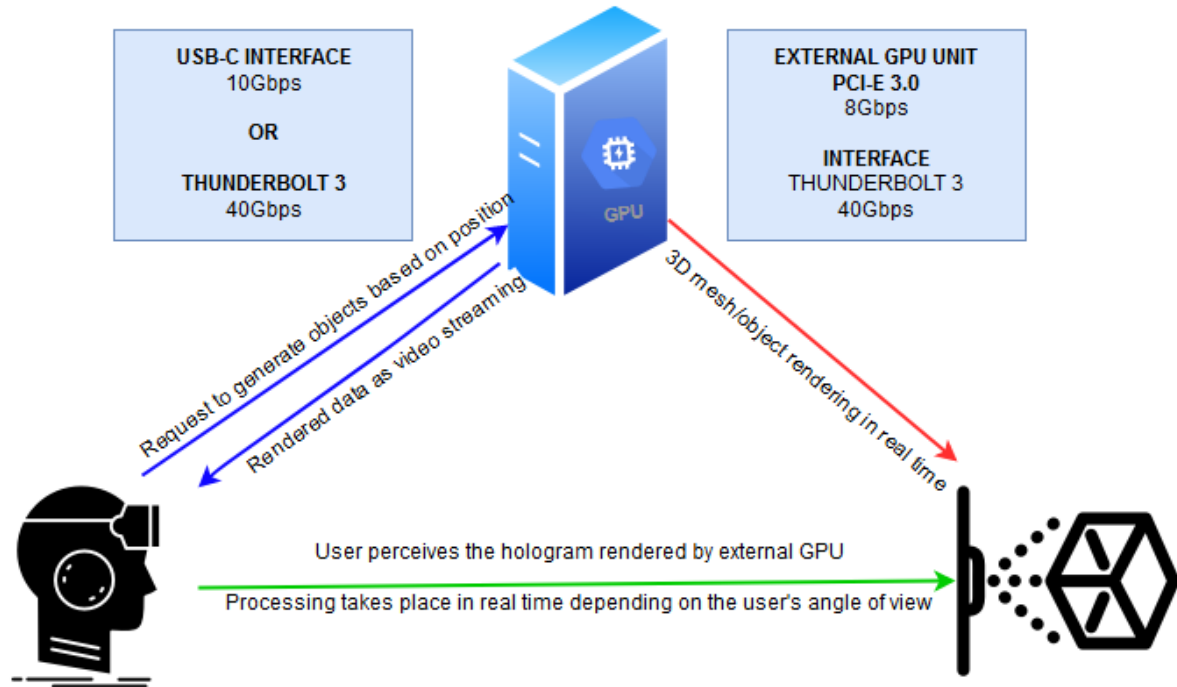


Figure 38: Diagram of how an external GPU solution would look like

As both USB-C and Thunderbolt 3 interfaces offer usually more theoretical bandwidth than the standard connection for GPUs (PCI-E 3.0) we should have no problem receiving all the processed data. However we must consider that usually the interface/connector specification sets a theoretical maximum limit, and it takes years and several iterations or versions of the same interface/connector to reach that outcome, so it is possible that in the case of USB-C the data transmission rate ends up being lower than the specification.

As we can see in the diagram above, the idea would be having the external GPU rendering in real time the LSIH to be displayed so that the HoloLens' HPU doesn't act as a bottleneck. Besides GPU power and having enough data transfer bandwidth so as for the HoloLens to receive the hologram's streaming, we would need to factor in latency as a key concept for this solution to be realistic.

This is definitely on-par with Microsoft's idea on this kind of solutions, as they plan to launch their Azure Kinect DX. This product is a developer kit that contains a best-in-class 1MP depth camera, 360° microphone array, 12MP RGB camera, and orientation sensor for building advanced computer vision and speech models. The idea here is to leave all of the spatial mapping processing to a device other than the HoloLens, and even though this device does not provide more CPU/GPU processing power, it does relieve some of the Spatial Mapping tasks the device would have to perform and thus liberating more power to be available to other tasks.

Second option: cloud computing/rendering solutions

As we mentioned before, Microsoft has promised that HoloLens 2 are being designed to work with their Azure Cloud Suite and to be able to access a Cloud Remote Rendering technology to boost the HoloLens' image processing capabilities. As we still don't know the specific details of this particular technology, we will outline how this could play out.

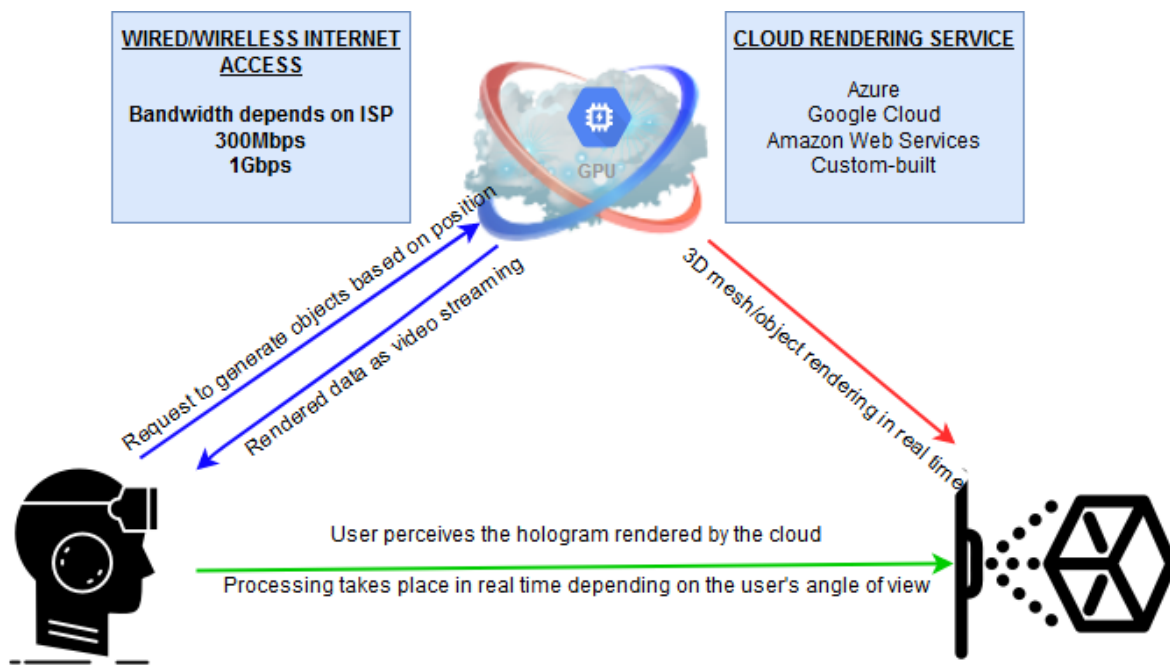


Figure 39: Diagram of a potential cloud rendering solution

The main principle of work remains similar to the one described in the previous solution. Again, we must rely in an external source of graphic computation to render the holograms to be seen in the device. However, this time instead of using an external GPU in a wired environment, we wouldn't even require additional equipment as we would be using a cloud solution.

This means that somewhere where the provider has its servers and datacenters, there will be a machine or series of machines in charge of rendering on-demand requests from this type of devices. Since this would be a cloud based service, we can draw a list of pros and cons that this approach has.

Pros	Cons
On-demand/Time-based service meaning that service does not need to be dedicated	Requires an internet connection with a good amount of bandwidth. No offline mode.
No need for the customer to require additional hardware or space to increase productivity	Latency/Delay must be incredibly low to deliver a proper experience
Added mobility for the subscribing devices	Data privacy might be compromised
Large variety of cloud-based hardware available to suit all types of needs	
Highly scalable, future-proof	

Table 3: Pros and Cons of rendering through Cloud Computing solutions

As we can see this approach has its downsides and its upsides. We will try to analyze whether or not these upsides surpass the downsides.

From an economic perspective, it is usually better to let others build the hardware infrastructure and features than to build it yourself. As noted in the first pro, with this kind of technology being used just in an on-demand basis, this means that:

1. From the user's point of view, we just need to pay for this service whenever we are making use of it. Rates could apply hourly, by the minute, or even depending on the computational power required to perform the tasks requested. This often results in a reduced bill depending on how intensive and frequent the need to use this service is. No additional hardware costs are needed.
2. From the service provider's point of view, the hardware infrastructure built to offer this type of service can be virtualized so as to allow multiple users to access it (simultaneously or not, depending on workload). Therefore, ensured that the service has enough customers and demand, we could potentially be exploiting this infrastructure to its maximum capacity, resulting in better payoff rates.

Secondly, we must note that by not having to invest in hardware (and its associated costs such as maintenance, etc.) as we did with the first solution and thus not having to upgrade it we do have a much easier opt-in opt-out policy, meaning that if the solution is not enough or ends up not working well, we would not have wasted a huge amount of money in trying to figure out its performance.

This point brings us to our third pro, that by not being bound to a specific external physical hardware, all of our devices making use of these cloud rendering services gain countless amount of mobility since we would just need to move the devices themselves. Since these devices are

typically HMD, they are very easy to transport. This could be especially meaningful for demonstrations for the sales team of a business (i.e. in the manufacturing section, an AR-based company trying to sell a fabric layout technology, etc.) and for congress, events and such.

The two last pros are kind of bound together. If this service has enough demand it will happen what has happened with other cloud services, the hardware specs and features that a cloud rendering service could offer will be within a very wide range of low-tier to high-end services and there is a certainty that every user would find a solution best fitted for their needs. With every hardware upgrade wave for GPUs or other hardware elements related to these services it will be the provider the one investing in it. So therefore there is almost no risk of these services granting outdated or sub-par services since there will most likely be enough competence in the market for the providers to try to not fall apart (Google Cloud, Azure, AWS...).

Now for the cons, we will see if there is a way for them to tone down or get a way around them. For the first one it does not seem like we have much of a solution currently. For accessing cloud-based services you will always need an internet connection. It is the base and principle of their existence, and that's why they are called 'cloud services'. So there is no turn around here: cloud rendering will require an internet connection. An offline mode is highly unlikely in the near future.

However, the real con here would be not just the requirement of an internet connection but which kind of internet connection. Rendering holograms designed to be classified as LSIH is not an easy task from a computational point of view. This will very likely mean that even though that

computational power is provided by the cloud rendering services, the data/video streaming that the device must receive in order to display LSIH would require a wide bandwidth to work with. This could be a limiting factor depending on the type of connection we have. With a traditional wired/wireless LAN connection we should have no problem as long as the ISP provides enough bandwidth. However, if using these devices in a more mobile setup we could end up setting up a LTE modem or tethering device to handle the connection, and there's where things could not go as expected.

The second con is closely related to the first one but does not only rely on internet and internet speed, but rather on the delay or latency that the connection has. We must remember that holograms are displayed based in their angle of view, so every movement from the user must be computed and taken into account to process the visible holograms. We could take two approaches here: either we send a raw, just processed data stream to the device which then handles that data locally (i.e. we send the whole building already rendered so that the device can display its various angles at a local scope) or we just send the data as a video stream. The first approach would be ideal to deal with delay but would still limit the amount and quality of the holograms to be displayed by the hardware that the device packs. The second one would be the desirable approach as long as the delay/latency remains consistent and low so that it does not generate a laggy experience.

For the last one that would mainly depend on our stance with these data privacy kind of issues. Handling our data with third parties is always riskier than processing them in a closed and proprietary environment as we would have if we were to pursue other solution. How this con

affects the overall solution depends on how sensitive the data to be sent/processed is, on what the policy of the service provider is to do with the data they receive...

Overall, we can state that cloud-based rendering solutions are very good suited to fulfill the requirements for deploying LSIH in devices such as the HoloLens and that although there are some cons to the approach, it remains as one of the best solutions nowadays.

Cloud solutions in the market

Although there are no current open-market real-time cloud rendering solutions we can see that there is a trend to bring them to the market as soon as possible. Microsoft has promised a service of this type for their new HoloLens v2 and there are other companies like Google investing in similar technologies (although with a different target in mind) as we could see in their Project Stream launch, now transitioned into Google Stadia.

However there are several companies offering cloud rendering solutions currently, although not tailored for real-time use but instead to render high density models, complex structures, and models requiring high computational power in general. These companies are highly likely to jump into the real-time cloud rendering trend if the demand is high enough as the adaptations they'd probably have to make are not that massive. Some of these companies are:

Google Cloud Rendering (Zync Render)

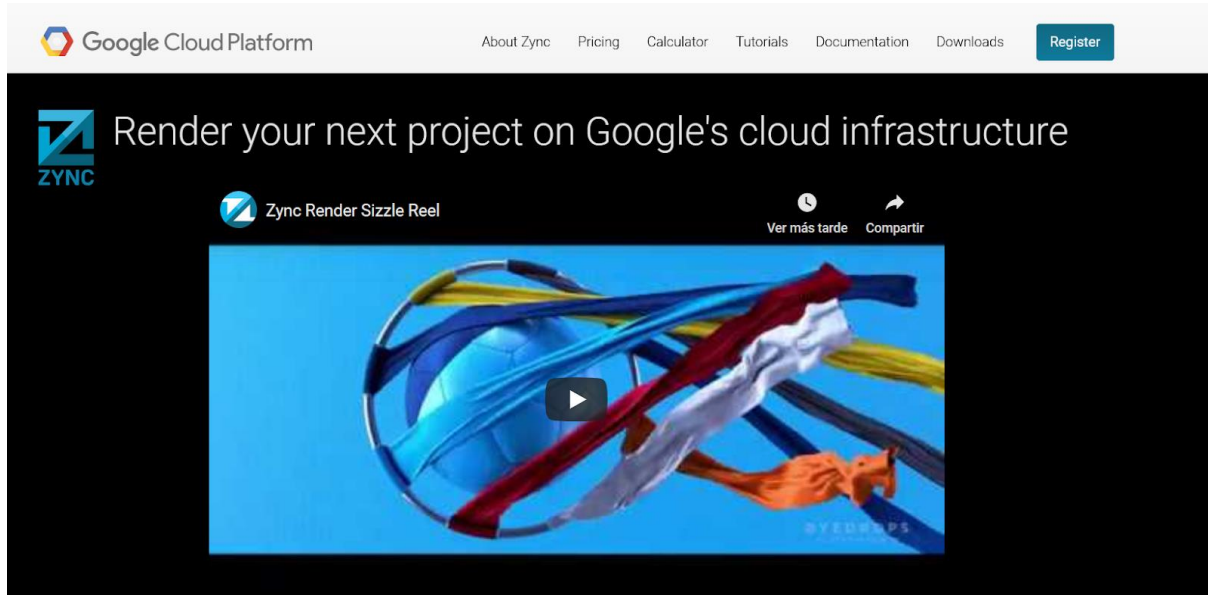


Figure 40: <https://www.zyncrender.com/>

Zync Render is a company that offers cloud rendering services by using Google cloud’s infrastructure to do so. They offer a wide variety of machines to choose from and offer access to several 3D modelling software suites. Therefore, the pricing depends on both the machine selected and the licensing for the software suite to be used as well.

Pricing

Software (costs are per hour):

Maya

Machine Type	V-Ray		Arnold		Renderman	
	Preemptible	Standard	Preemptible	Standard	Preemptible	Standard
zync-8vcpu-16gb	\$0.51	\$0.79	\$0.98	\$1.26	\$0.56	\$0.84
zync-16vcpu-32gb	\$0.65	\$1.19	\$1.05	\$1.59	\$0.63	\$1.17
zync-32vcpu-64gb	\$0.93	\$2.01	\$1.19	\$2.27	\$0.77	\$1.85
zync-64vcpu-128gb	\$1.51	\$3.19	\$1.48	\$3.61	\$1.06	\$3.19
zync-96vcpu-128gb	\$2.02	\$4.99	\$1.70	\$4.67	\$1.28	\$4.25
zync-96vcpu-192gb	\$2.08	\$5.28	\$1.76	\$4.96	\$1.34	\$4.54
zync-8vcpu-64gb-1gpu-p100	\$1.50	\$3.02				
zync-8vcpu-64gb-4gpu-p100	\$4.29	\$8.90				

Figure 41: Pricing for Zync Render if using Maya. Data storage is charged separately

Amazon Sumerian:

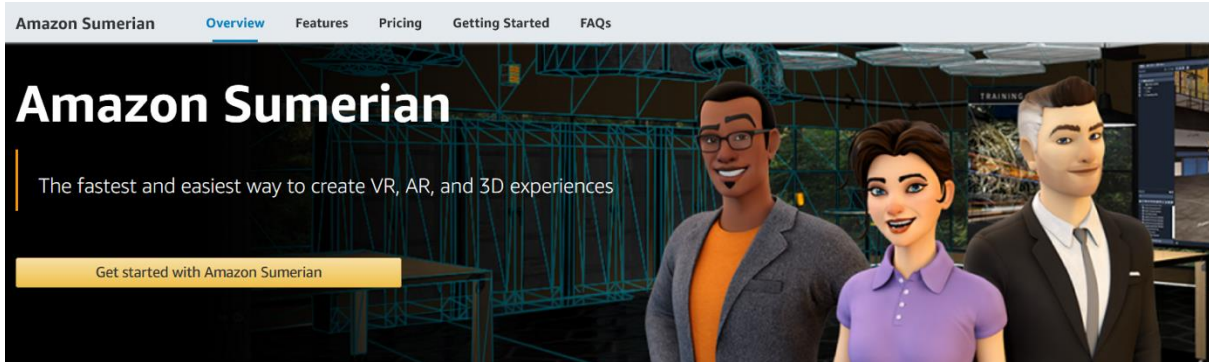


Figure 42: Amazon Sumerian portal from their website

Here we can appreciate that there's another of the big companies that has jumped straight into these types of solutions. Although the solution here is a little bit different from the others. Instead of offering cloud rendering solutions per se, they provide an environment, similar to Unity, in which users can build and deploy (cloud-based) VR or AR applications.

For the pricing, as it is a different solution that the other ones, the system is also slightly different. They charge you for the amount of stored data as well as the traffic that that data generates.

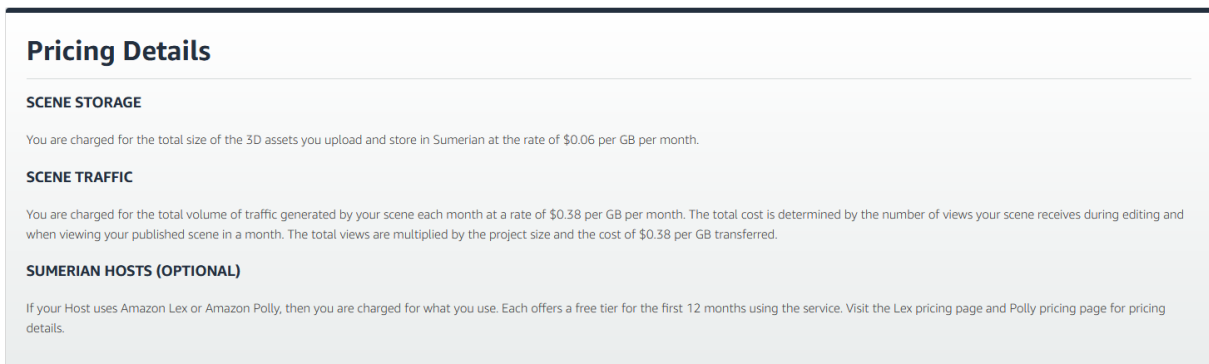


Figure 43: Pricing details for Amazon Sumerian service

GarageFarm / Xesktop:




Figure 44: Xesktop server specifications from their website

Contrary to the previous described service, Xesktop only offers the advertised servers as a whole package to be rented for a given amount of time. However, they do offer discounts based on how much you pay in advance, as can be seen in Figure 45. These servers are rented by hours.

amount paid	free %	free credits	total credits	hours	1server/hour
\$250 USD	4%	\$10	\$260	43.33	\$5.77
\$500 USD	4%	\$20	\$520	86.67	\$5.77
\$1,000 USD	8%	\$80	\$1,080	180.00	\$5.56
\$1,500 USD	12%	\$180	\$1,680	280.00	\$5.36
\$2,000 USD	16%	\$320	\$2,320	386.67	\$5.17
\$2,500 USD	20%	\$500	\$3,000	500.00	\$5.00
\$3,000 USD	24%	\$720	\$3,720	620.00	\$4.84
\$3,500 USD	28%	\$980	\$4,480	746.67	\$4.69
\$4,000 USD	32%	\$1,280	\$5,280	880.00	\$4.55
\$4,500 USD	36%	\$1,620	\$6,120	1,020.00	\$4.41
\$5,000 USD	40%	\$2,000	\$7,000	1,166.67	\$4.29
\$5,500 USD	44%	\$2,420 USD	\$7,920 USD	1,320.00	\$4.17 USD
\$6,000 USD	48%	\$2,880 USD	\$8,880 USD	1,480.00	\$4.05 USD
\$6,500 USD	52%	\$3,380 USD	\$9,880 USD	1,646.67	\$3.95 USD

Figure 45: Pricing for Xesktop based on how much payment is done in advance

Fox Renderfarm:

Figure 46: Costs for a high-end service at FoxRenderfarm

Rendershot:

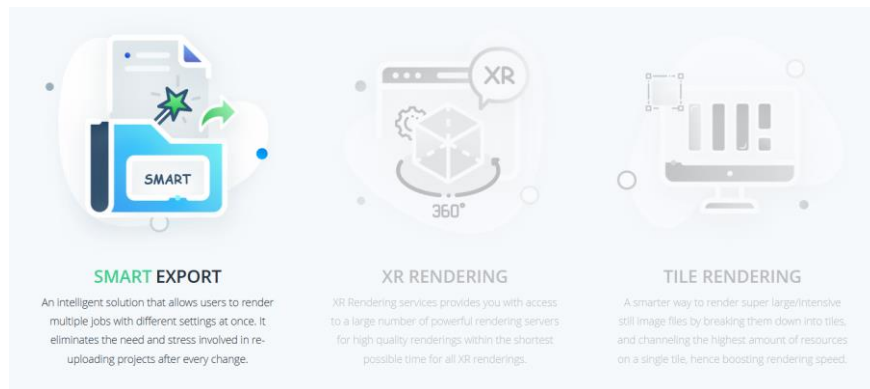


Figure 47: Some of rendershot's features

Rendershot is currently one of the few companies that list XR ('X' mentioning the possibility of the X being substituted by V, A or M for VR/MR/AR) Rendering as a future feature supported by their servers.

Pricing is based on computational power and also depends on the server specs being used. A priority system is implemented for traditional rendering requests so it gets cheaper the longer you are willing to wait in a queue.

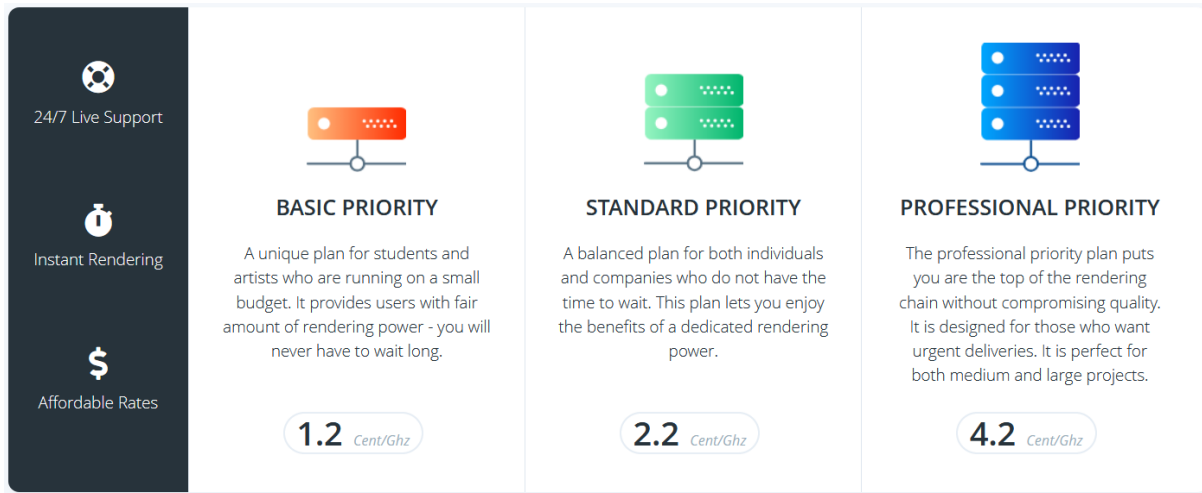


Figure 48: Pricing for Rendershot services

So we can conclude that although current cloud rendering solutions do exist, there is still nothing in the public market that allows to render models in real time. However, there are plenty of clues that can almost guarantee that this is not an utopic scenario and that companies are investing in this kind of technology. And given that a big company like Google is investing too, other companies are very likely to follow the lead.

IX. Conclusions

Throughout this project we have set up a series of goals and objectives to which we will now refer as our guideline to draw the conclusions we can extract from this document.

The first main objective we set up was to determine whether or not LIDAR technology is suitable to generate Large Scale Immersive Hologram 3D models. As we have seen throughout the correspondent methodology section, LIDAR scans offer a very detailed reconstruction of the scanned area. These results are better if we overlap several scans performed over the same area but from different angles. We have covered that there are currently several techniques and methods to render those generated point-clouds and to make them into a readable mesh object.

We have also seen how these mesh objects can be further processed and suited to whatever environment we will be using them on. So therefore we can come to the conclusion that LIDAR scanning technology is good enough so as to provide 3D scenarios or models from where we could extract LSIH objects.

However, one of the main objectives for this project, which was actually deploying those LSIH 3D objects in the HoloLens device did not go exactly as planned. Mainly due to the technical and hardware-level limitations that the device packs, the deployment of these LIDAR-scanned based models yielded a poor performance unless serious quality reductions were made. This did not imply that the models couldn't run in the device, but rather that the device still is not powerful enough so as to host higher quality versions of that model. Besides, the limitations do not factor only the computational power of the device but also the limited MR field of view that their visor

is able to offer. The region in which a user can see the holograms being displayed is very limited when compared to the whole range of vision a human eye has, losing the ‘immersive’ aspect of the interaction. Thus, for this point we can conclude that LSIH deployed in HoloLens, at least for more complex models such as buildings, still have a great room for improvement.

This sets up the stage for the last part of the document, where we outline what is very likely to happen in the near future with these technologies. Microsoft has already announced its next iteration of the HoloLens device with significantly more processing power and an amplified field of vision which could in turn result in a more immersive experience. This, together with the growth of AR and MR technologies in general now that people begin to realize the number of applications those technologies have, set up a trend for Microsoft and other companies to continue investing in these devices.

Even though the new devices will pack more power in their SoC, we have to remember that they will remain as mobile HMD devices, and therefore their hardware limitations will still drag the overall capabilities down. That is why we outline that using external CPU/GPU processing, either via traditional cable interfaces (USB-C, Thunderbolt) or Cloud Computing technologies, will greatly contribute in the future towards the achievement of a realistic LSIH experience within these devices.

X. Bibliography

- [1] Nee, A. Y., Ong, S. K., Chryssolouris, G., & Mourtzis, D. (2012). *Augmented reality applications in design and manufacturing*. CIRP Annals-manufacturing technology, 657-679.
- [2] Dong S, Kamat VR (2010) *Robust Mobile Computing Framework for Visualization of Simulated Processes in Augmented Reality*. Proceedings of the 2010 Winter Simulation Conference (WSC'10), 3111–3122.
- [3] Kress, B. C., & Cummings, W. J. (2017, May). 11 - 1: Invited Paper: *Towards the Ultimate Mixed Reality Experience: HoloLens Display Architecture Choices*. In SID Symposium Digest of Technical Papers (Vol. 48, No. 1, pp. 127-131).
- [4] Hockett, P., & Ingleby, T. (2016). *Augmented reality with HoloLens: Experiential architectures embedded in the real world*. arXiv preprint arXiv:1610.04281.
- [5] Baker, N. “*Mixed Reality*” *Keynote at Hot Chips HC28 – Symposium for High Performance Chips*, Aug. 21-23 2016, (www.hotchips.org).
- [6] Faggiano, A. *Study and prospects of large-scaled immersive holograms using the Microsoft HoloLens*. IIT Graduate Research Paper.
- [7] Vicent, L. *Large Scale Immersive Holograms for prevention and education purposes*. IIT Graduate Research Paper.
- [8] Arguillère, P. *Telepresence Framework – 3D Interfaces, Large Scale Immersive Holograms*. IIT Graduate Research Paper.
- [9] Al-Habash, N. *Large scaled immersive holograms for prevention*. IIT Graduate Research Paper.

- [10] Benchmark comparison for Intel Atom X5 and Qualcomm Snapdragon 845:
<https://versus.com/en/intel-atom-x5-z8500-vs-qualcomm-snapdragon-845>
- [11] Fraiss, S.M. *Rendering Large Cloud Points in Unity*. BSc Thesis, Technischen Universität Wien
- [12] Choi, J., & Ko, J. (2019, June). *RemoteGL-Towards Low-latency Interactive Cloud Graphics Experience for Mobile Devices*. In Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services (pp. 693-694). ACM.
- [13] Krol, Z. M. T. *Cloud Gaming: Implications of cloud gaming, shifting towards servitization in the gaming industry*.
- [14] Harris, M. J., & Lastra, A. (2001, September). *Real - time cloud rendering*. In Computer Graphics Forum (Vol. 20, No. 3, pp. 76-85). Oxford, UK and Boston, USA: Blackwell Publishers Ltd.
- [15] Harris, M. J., & Lastra, A. (2002, March). *Real-time cloud rendering for games*. In Proceedings of Game Developers Conference (pp. 21-29).
- [16] Elek, O., Ritschel, T., Wilkie, A., & Seidel, H. P. (2012). *Interactive cloud rendering using temporally coherent photon mapping*. Computers & Graphics, 36(8), 1109-1118.
- [17] Hoover, Melynda, *An evaluation of the Microsoft HoloLens for a manufacturing-guided assembly task* (2018). Graduate Theses and Dissertations. 16378.