

# Exploring Extended Reality with ILLIXR: A New Playground for Architecture Research

Muhammad Huzaifa Rishi Desai Xutao Jiang Joseph Ravichandran Finn Sinclair Sarita V. Adve  
University of Illinois at Urbana-Champaign  
illixr@cs.illinois.edu

**Abstract**—As the need for specialization increases and architectures become increasingly domain-specific, it is important for architects to understand the requirements of emerging application domains. Augmented and virtual reality (AR/VR) or extended reality (XR) is one such important domain. This paper presents a generic XR workflow and the first benchmark suite, ILLIXR (Illinois Extended Reality Benchmark Suite), that represents key computations from this workflow. Our analysis shows a large number of interesting implications for architects, including demanding performance, thermal, and energy requirements and a large diversity of critical tasks such that an accelerator per task is likely to overshoot area constraints. ILLIXR and our analysis have the potential to propel new directions in architecture research in general, and impact XR in particular.

## I. INTRODUCTION

Recent years have seen the convergence of multiple disruptive trends to fundamentally change computer architecture: (1) With the end of Dennard scaling and Moore’s law, application specialization has emerged as a key architectural technique to meet the requirements of emerging applications, (2) computing and data availability have reached an inflection point that is enabling a number of new application domains, and (3) these applications are increasingly deployed on resource-constrained edge devices, where they interface directly with the end-user and the physical world. In response to these trends, our research conferences have seen an explosion of papers on highly efficient accelerators, many of which are focused on machine learning. Thus, today’s computer architecture researchers must not only be familiar with hardware principles, but more than ever before, must understand emerging applications. To truly achieve the promise of efficient edge computing, however, will require architects to broaden their portfolio from specialization for individual accelerators to understanding domain-specific *systems* which may consist of multiple sub-domains requiring multiple accelerators that interact with each other and potentially the cloud to collectively meet the end-user demands for that domain.

This paper makes the case that the emerging domain of augmented, virtual, and mixed reality, collectively referred to as extended reality (XR), is a rich domain that can propel architecture research on efficient domain-specific edge systems. Our case rests on the following observations: (1) XR will pervade most aspects of our lives – it will affect the way we teach, conduct science, practice medicine, entertain ourselves, train professionals, our social interactions, and more – XR is envisioned to be the next interface for most of computing [47] [16]. (2) While current XR systems exist today, they have a long way to go to provide a tetherless experience approaching perceptual abilities of humans. As discussed later, there

is a gap of several orders of magnitude in performance, power, and usability, giving architects a potentially rich space to optimize. (3) XR involves a large number of diverse sub-domains – audio, video, graphics, haptics, optics, and robotics – making it challenging to design a system that executes each of the above well while fitting in the resource constraints. (4) The combination of realtime constraints, a complex pipeline consisting of interacting components, and ever-changing algorithms provides an opportunity for full stack optimizations involving the hardware, runtime and OS, programming system, and the application. Given that the end-consumer is a human with limited perception enables the full stack optimization to consider approximation as a way to achieve the constraints.

A key obstacle to doing architecture research for XR is that there are no benchmark suites covering the entire XR workflow and constituent sub-domains to drive such a research. This paper develops and characterizes ILLIXR<sup>1</sup> (Illinois Extended Reality Benchmark Suite), the first open-source benchmark suite to represent the XR domain. We developed this suite after consultation with many colleagues and researchers in the sub-domains mentioned above. These consultations led to identifying a representative, generic XR workflow. We then identified key components within this workflow. We collected and modified open source codes for many of these components which constitute our benchmark suite. Finally, we analyze the performance and energy characteristics of this suite on a desktop class and an embedded class machine and draw several implications for future architecture research.

Specifically, this paper makes the following contributions:

- We describe a generic XR workflow and key representative components. We develop ILLIXR, the first open source benchmark suite for XR research in architecture, containing representative components from the above workflow.
- We report performance and power on a desktop and embedded machine for each component. We find they are collectively far from the power and performance requirements for future devices, making the case for efficiency through specialization.
- Our detailed analysis shows that the components constituting a modern XR workflow are quite diverse with new algorithms continuously evolving, making the case for flexible, programmable hardware.
- Our detailed analysis shows that the components are quite diverse in their requirements with respect to each other as well as consist of diverse tasks within themselves, with no clear

<sup>1</sup>Pronounced elixir.

TABLE I: Ideal requirements of AR and VR vs. state-of-the-art devices, HTC Vive Pro for VR and Microsoft HoloLens 2 for AR. Silicon area and power for HTC Vive Pro are assumed to be same as typical die sizes and TDPs of desktop CPUs and GPUs. VR devices are typically larger and so afford more power and thermal headroom. Our ideal case requirements are based on sizes of SoCs found in a VR headset, such as Snapdragon 835 in Oculus Quest, and small AR glasses, such as APQ809w in North Focals.

Metric	HTC Vive Pro [6]	Ideal VR [25], [35]	Microsoft HoloLens 2 [13]	Ideal AR [25], [35]
Resolution (MPixels)	4.6 [6]	200	4.4	200
Field-of-view (Degrees)	110 [6]	Full: 165×175 Stereo: 120×135	52 diagonal [31], [72]	Full: 165×175 Stereo: 120×135
Refresh rate (Hz)	90 [6]	90 – 144	120 [63]	90 – 144
Motion-to-photon latency (ms)	< 20 [37]	< 20	< 2 [76]	< 5
Power (W)	–	1 - 2	10 [15], [72]	0.1 – 0.2
Silicon area ( $mm^2$ )	–	100 – 200	79 [15], [72]	10s
Weight (grams)	470 [73]	100 – 200	566 [13]	10s

dominant tasks emerging. This diversity poses a challenge to specialization – an accelerator per task will overshoot area and other requirements. Thus, this analysis motivates shared hardware with “generalizable” specialization that can be availed by multiple tasks.

We note that our goal is not to make a commercial grade XR system, either in terms of its completeness or in terms of its performance and efficiency. Our goal instead is to provide researchers with a suite consisting of components representative of an XR workflow based on the state-of-the-art algorithms for architects and system designers to experiment with. Some missing attributes of ILLIXR that we leave for future work are: we do not yet have codes for some parts of the workflow; we do not model OS or real-time schedulers or work partitioning with the cloud; and we currently analyze the components mostly in isolation.

ILLIXR is open-source and available at <https://illixr.github.io>. Our hope is that others will contribute to make ILLIXR a growing repository of an open source XR benchmark suite that can drive future architecture and system design in emerging domain-specific edge devices.

## II. XR SCENARIOS AND REQUIREMENTS

### A. Definitions

Virtual reality (VR) immerses the user in a completely digital environment with realistic sensory input. VR is defined by the use of opaque displays and immersive content. In contrast, augmented Reality (AR) enhances our real world with overlaid contextual content. Mixed reality (MR) is often used as a middle-ground between fully opaque experiences and contextual AR. Extended Reality (XR) is being adopted as a general umbrella term for AR, VR, and MR. Consequently, throughout the rest of this paper, we use XR to refer to all technologies collectively.

### B. Requirements

XR requires devices that are lightweight, mobile, and all day wearable, and yet provide enough compute performance and low thermals and energy consumption to meet the demands of the intensive applications they are meant to enable. The traditional PPA (power-performance-area) area metric is therefore important, but not a complete characterization. Performance characteristics include

resolution, field-of-view (FoV), and refresh rate of the display, and the end-to-end or motion-to-photon latency of the system. These metrics interact with each other; e.g., higher resolution, FoV, and refresh rate can adversely affect motion to photon latency. In addition to PPA, size, weight, area, and power of the devices, usually abbreviated as SWAP, are equally important metrics for determining the usability of the device. We summarize these metrics in Table I.

Table I summarizes these requirements for current state-of-the-art VR and AR devices and the aspiration for ideal futuristic devices. These requirements are driven by both human anatomy and usage considerations: ultra high performance to mimic the human visual system; extreme thermal constraints due to contact with human skin; super lightweight due to comfort; and small form factor due to social acceptability and ease of use. We identified the values of the various aspirational metrics through an extensive survey of the literature [35], [47], [48]. Although there is no consensus on exact figures, there is an evolving consensus on approximate values that shows orders of magnitude difference between the requirements of future devices and what is implemented today (making the exact values less important for our purpose). For example, the power gap alone is two orders of magnitude. Coupled with the other gaps (e.g., two orders of magnitude for resolution), the overall system gap is many orders of magnitude across all metrics.

The implications of this gap are enormous. While architects often deal with power and performance gaps, the rest of the SWAP gap also severely limits the number of sensors – cameras (for SLAM, eye tracking, video recording, and hand tracking), IMUs, microphones, speakers, LTE and WiFi. Consequently, the system and algorithms need to be adapted to work in such a constrained environment, and it is likely that such a large gap can only be eliminated via synergistic advances throughout the computing stack: algorithms, applications, runtimes, and hardware. There has been significant progress in algorithmic techniques to bridge some of these large gaps. For example, foveated rendering partly addresses the gap in display resolution by leveraging the anatomy of the eye and certain aspects of the human visual system. Time warp is now a standard component that enables acceptable motion to photon latency by applying fast predictors to enhance the slower conventional tracking algorithms. While it does not affect the power usage of the tracking algorithm, good predictions allow slower but accurate tracking to be invoked at lower rates, thereby saving power. These continuing algorithmic advances and their interactions among the different components reinforce the need for programmable hardware specialization that is codesigned with other layers in the system stack. This paper provides state-of-the-art implementations of key components of the XR pipeline to enable architects and system designers to participate in such a codesign process.

## III. A GENERIC XR WORKFLOW

This section describes a generic XR workflow, summarized in Figure 1. The workflow describes the computation that a generic XR device must perform from the moment the user provides an input to the moment the display and audio are updated. We distilled this workflow from multiple sources that describe different parts of a VR, AR, or MR pipeline. For completeness, our workflow assumes all computation is done at the device (i.e., no offloading to a server or the cloud). For simplicity, we assume a single user;

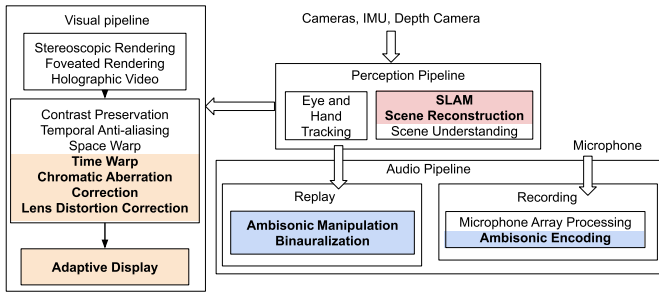


Fig. 1: A generic XR workflow. The components provided by ILLIXR are colored.

i.e., there is no communication with multiple XR devices. Thus, our workflow does not represent any network communication. We leave the integration of networking, edge and cloud work partitioning, and multiparty XR to future work. Figure 1 also shows the components of the workflow that are provided by ILLIXR in shaded boxes, as discussed in more detail in Section IV.

The input to the XR system comes in the form of data collected by multiple sensors. We generalize the required input as some form of the user’s motion – head rotation, eye movement, and hand movement – and speech. To capture motion, we assume one or more cameras and one or more IMUs (Inertial Measurement Units, typically consisting of an accelerometer and gyroscope). The combination of camera input (infrequent, but less noisy) and IMU input (frequent, but more noisy) is used to improve both accuracy and latency of the system as described below. For speech, we assume microphone arrays.

The sensor inputs feed into three interacting pipelines: (1) the *perception* pipeline, which develops an understanding of where the user is and what is being viewed (for AR/MR) in a form usable by the rest of the system; (2) the *visual* pipeline, which takes the input from the perception pipeline to render and display the next frame as seen from the user’s new vantage point; and (3) the *audio* pipeline which also takes the perception pipeline’s information to encode/play 3D spatial audio, given the new position of the user. We discuss each of these three pipelines next.

### A. Perception pipeline

The primary goal of the perception pipeline is to take the user’s physical motion related inputs from the sensors and translate them into information understandable to the rest of the system so it can render and play the new scene and sound for the user’s new position.

Broadly, this stage includes components such as head tracking (the orientation of the user’s head, performed typically using a SLAM algorithm (see below)), eye tracking (the gaze of the user’s eye), hand tracking (gesture recognition), scene depth estimation (how far objects in the user’s surroundings are), and, if needed, scene reconstruction and understanding (for MR and AR).

**ILLIXR components.** ILLIXR focuses on two of the most important components in the perception pipeline. The first, head tracking, is achieved through Simultaneous Localization and Mapping (SLAM), a technique that can track the user’s motion without the use of external markers in the environment (Section IV-A1). SLAM determines where the user is in the world and what the world looks like. For the purposes of head tracking, we

use a sparse algorithm with its output conveyed to both the visual and audio pipelines through a six degrees of freedom pose (6 DOF pose).

The second component of the perception pipeline supported in ILLIXR is scene reconstruction (Section IV-A2). This is needed to enable MR and perform realistic physics in AR and uses a dense SLAM algorithm optimized to enable scene reconstruction.

### B. Visual pipeline

The visual pipeline takes information pertaining to the user’s new position from the perception pipeline and produces the final display. We divide this process into three parts: rendering, post-processing, and display, as described below.

**Rendering.** The rendering part of the visual pipeline takes the outputs from the perception pipeline to determine how the displayed scene should be changed and renders the updated scene. This part draws significantly from conventional graphics technology, but needs to incorporate emerging algorithms to account for the far more demanding visual requirements of XR (e.g., larger resolution and field of view). For example, foveated rendering exploits the fact that the eye uses the highest fidelity only in the part of the image close to the fovea; the peripheral image is acceptable with lower fidelity. This algorithmic innovation enables significant reduction of computation from the standard constant-resolution rendering techniques.

**Post-processing.** This part processes the rendered images to correct for imperfections due to both the rendering process and in the display optics. While the individual computations in this part are not as expensive as the rendering part, they contribute significantly to the user experience.

The first set of computations include contrast preservation and temporal anti-aliasing to improve contrast in the periphery of the image and to remove flickering of pixels from frame to frame, respectively (both artifacts are common in traditional rendering but become perceptible in XR).

The second set of computations takes the image generated thus far and modifies it to account for the user’s movements that happened between the time the rendering process started and the time the rendering process finished. Instead of re-rendering the image upon any head and eye movement, this computation, called time warp, takes the available rendered image and simply changes the perspective to account for (a prediction of) the user’s movement; i.e., this is an optimization that improves the user experience by not showing older, stale frames with the wrong perspective. The presence of the IMU enables this computation by providing pose changes far more frequently (albeit with possibly more noise) than the SLAM algorithm used in the perception pipeline. Using the noisy IMU data typically suffices for this level of correction (but is insufficient as the only mechanism to track the user’s pose). This use of the IMU coupled with time warp is key to meeting the required motion to photon latency – by predicting the user’s movement for short bursts, it is possible for the visual pipeline to run asynchronously relative to the perception pipeline and is not held back due to the high computational demand of that pipeline. (Similarly, timewarp is also used to compensate for frames that might otherwise be dropped because the renderer could not finish on time.) Timewarp is critical in XR applications, as any perceived latency will cause discomfort [26], [49], [70].



The third set of computations are lens distortion and chromatic aberration corrections. Most VR devices use small displays placed close to the user’s eyes; therefore, some optical solution is required to create a comfortable focal distance for the user [36]. These optics come with the consequence of large amounts of optical and chromatic distortion – lines become curved and colors get split into the RGB spectrum, respectively. These optical affects are corrected by applying *reverse* distortions to the rendered image, such that the distortion from the optics is cancelled out [68].

**Display.** The final stage of the visual pipeline updates the display. The display computations take the final image after post-processing and modify it for consumption by the optical devices. While modern display optics are simply a combination of fixed lenses and LCDs, future displays are anticipated to be adaptive and multi-focal, capable of generating true 3D imagery using computational holography. Such displays require significant computation to map the rendered image to the optical system (Section IV-B2).

**ILLIXR components.** For the visual pipeline, ILLIXR supports timewarp, lens distortion correction, chromatic aberration correction, and computational holography for the display. Most implementations integrate timewarp, lens distortion, and chromatic aberration correction; we therefore provide a single component for these and refer to it collectively as *timewarp*. We refer to the display component as *hologram*.

### C. Audio pipeline

The audio pipeline is responsible for generating realistic spatial audio and can be divided into recording and playback. During recording, microphone array processing and Ambisonic encoding create an Ambisonic soundfield that can be later played back. Microphone array processing involves combining signals from various microphones and Ambisonic encoding consists of combining all the physical source signals into an Ambisonic soundfield.

During playback, Ambisonic soundfield manipulation is performed using SLAM information. This includes rotation and zooming of the soundfield as the user’s pose changes. Once the final soundfield is obtained, it is mapped to physical speakers, typically headphones in XR devices. This process is called binauralization. It adds certain cues to the sound that are present in real life; e.g., the modification of incoming sound by the person’s nose, head, and outer ear shape. Such cues allow the digitally rendered sound to mimic real sound, and thus help the brain localize sounds akin to real life.

**ILLIXR components.** Within the audio pipeline, ILLIXR supports Ambisonic encoding for recording, and Ambisonic manipulation and binauralization for playback.

### D. Putting it Together: A Moment in the Life of an XR User

To better illustrate the overall flow, we take the example of a simple MR application where a user is playing ping-pong against the computer in their living room, which has been transformed to be part of an Olympic stadium (this example was chosen because it exercises all three computing pipelines). The perception pipeline continuously processes all the incoming sensor data to estimate the hand position, eye position, and location of the user. The perception pipeline

TABLE II: Components currently in ILLIXR. OpenCV is a popular computer vision library and GLSL stands for OpenGL Shading Language.

Component	Algorithm	Implementation
<b>Perception Pipeline</b>		
SLAM (VIO)	OpenVINS [14]	C++, OpenCV
Scene Reconstruction	ElasticFusion [5]	C++, CUDA, GLSL
<b>Visual Pipeline</b>		
Time Warp	VP-matrix reprojection using pose [68]	C++, GLSL
Lens Distortion	Mesh-based radial distortion [68]	C++, GLSL
Chromatic Aberration Correction	Mesh-based radial distortion [68]	C++, GLSL
Adaptive display	Weighted GerchbergSaxton (GSW) hologram generation [56]	CUDA
<b>Audio Pipeline</b>		
Recording	Ambisonic encoding [12]	C++
Playback	Ambisonic manipulation [12]; binauralization [12]	C++

also updates the reconstruction of the scene using this data and passes all this information to the visual pipeline and audio pipelines.

The game engine uses this information for calculating the physics of the ball and handling occlusion, and then kicks off the visual pipeline to render the scene from the new perspective. Then the image in the framebuffer is warped to account for any user movement that may have happened between the time the rendering process began and by the time it ended. (The time warp code uses updated information from the IMU to get its new prediction of the user’s pose, since the IMU operates at a much faster rate than the camera which feeds SLAM.) Then, the image is corrected for lens and chromatic aberrations caused by the curved lenses of XR wearables. Finally, the pixels are modified using computational holography to map them to different depth planes to present a true 3D image to the user.

The audio pipeline uses the information from the input stage for correctly modeling acoustics (reflections, reverb, etc.), and for rotating and zooming the virtual soundfield to account for the new pose. It then maps the soundfield to the user’s headphones with correct cues (binauralization).

## IV. ILLIXR COMPONENTS

This section describes the components included in ILLIXR. Table II shows the name, algorithm, and implementation information of each of these components.

### A. Perception Pipeline

1) *SLAM*: There are a large number of distinct algorithms that have been proposed for SLAM for a variety of scenarios, including for robots, drones, VR, and AR. The constraints, requirements, and sensor inputs for different scenarios are different, leading to different optimal choices; e.g., some scenarios only require odometry (enabling sparse algorithms) while others require detailed scene reconstruction (favoring dense algorithms). In this work, for the purposes of pose estimation for XR, we use the OpenVINS algorithm and its open source implementation [14]. This algorithm uses two cameras and an IMU for sensor input for higher accuracy (referred to as visual-inertial odometry or VIO). OpenVINS has

been shown to be more accurate than other popular algorithms such as VINS-Mono, VINS-Fusion, and OKVIS [55].<sup>2</sup>

The frontend of OpenVINS is a tracker that detects new features in both eye images and matches them with other images. The feature detector first divides the incoming left eye image into a grid of  $8 \times 5$  pixels and extracts FAST features [60] from each grid cell in just this image. Then, it uses the KanadeLucasTomasi (KLT) [34], [42] feature tracker to project features from the left camera to the right to generate a list of features in both cameras that the feature matching stage can use. Feature matching includes both temporal matching – two different images in time for the same eye – and stereo matching – left and right eye images for the same moment in time. Matching is performed using KLT and then the matched features are passed through a RANSAC algorithm to remove outliers. The inliers are then added to the feature database.

After performing tracking, OpenVINS uses a Multi-State Constrained Kalman Filter (MSCKF) [51] to estimate the user’s current position (which could have changed since the camera input). This starts by processing the IMU data between the current frame and the previous frame (10 readings in our experiments) and estimating the final position of the user. Then OpenVINS uses the features saved by the tracker and the IMU-based position to update the estimation filter. The update process consists of: 1) a two-step triangulation of the features (a linear triangulation and a Gauss-Newton refining process), 2) a linearization of the MSCKF features (by computing the Jacobian, projecting the nullspace, and performing a Chi squared distance check), 3) Givens rotations to compress all of the measurements in the feature matrix, and 4) the actual update of the Kalman filter.

The open source version of OpenVINS was built on top of ROS (Robot Operating System) [3] as is common for most robotics software. Since ROS is known to incur non-trivial overheads, we removed the ROS dependence from the code and created a standalone version that is more appropriate for architects.

2) *Reconstruction*: For Scene Reconstruction, we chose ElasticFusion [71], a dense SLAM algorithm that uses an RGB-Depth (RGB-D) camera to build up a dense 3D map of the world. ElasticFusion first converts the incoming RGB-D images into textures and removes noise from the depth data with bilateral filtering. This is the camera processing stage of ElasticFusion.

In the image processing stage, ElasticFusion converts the filtered depth data into a vertex map where each depth point is converted to a point in 3D Cartesian space. It then uses the vertex map to compute a normal map. Data layout transformations are also performed in this stage, where RGB\_RGB is converted to RR\_GG\_BB [5].

Each point in the map (also called the model) is represented as a surfel which includes the position, normal, color, radius, and timestamp of the map point. The surfel prediction step involves predicting which surfels are active and which are not. The predicted surfel model is then used to perform the first phase of tracking [71].

Using both the vertex and normal maps, a transformation matrix is computed to estimate the new camera pose. This is a multi-step

process in which first ICP (Iterative Closest Point) is used to find similar points between the new image and the predicted surfel model (frame-to-model tracking), and then the corresponding geometric error is computed [71]. Then, photometric error between the image and the model is computed, and the total error is obtained by combining geometric and photometric errors.

Then, ElasticFusion attempts to perform global loop closure. Global loop closure implies that a previously visited location has been revisited, and the accumulated drift since then can be reset. If global loop closure is not possible, then the pose is estimated by comparing the new predicted surfel model and the previous frame’s surfel model (model-to-model tracking). This pose estimation is identical to the frame-to-model pose estimation. Finally, the computed pose is used to fuse the information from the new frame into the map [71].

## B. Visual Pipeline

1) *Timewarp, lens distortion, and chromatic aberration*: The timewarp, lens distortion, and chromatic aberration correction implementations are reverse engineered from [11]. We extracted the distortion and timewarp shaders from this reference, and built our own program using OpenGL FBOs<sup>3</sup> as a common interface between neighboring benchmark components. We integrated timewarp, lens distortion, and chromatic aberration into one component.

The reference implementation [11] performs lens distortion and chromatic aberration correction using three independent sets of UV coordinates, each corresponding to one color channel. The image to be processed is mapped onto a flat plane using distorted UV coordinates.

The distorted UV coordinates, calculated using the physical properties of the lens, result in a barrel distortion filter [68]. To calculate the chromatic aberration, each color channel’s scaling factor is multiplied by its angle from the center line, and the chromatic aberration coefficients of the lens itself.

The shader samples the eye texture three times for each pixel, once per color channel. This has implications for GPU texture cache locality, as the texture accesses for each individual color channel will not be congruent due to distortion.

Timewarp is performed by recording the user’s pose both at the beginning of the frame and just before vsync, and the viewproj matrix used for the rendering process [68]. First, the new view matrix is computed using the new pose. Then, the timewarp transformation matrix is computed using the projection matrix and the delta between the old and new view matrices. Finally, the transform is applied to each color channel’s UV coordinates and the rendered image is sampled at the new coordinates to obtain the warped image. Timewarp is performed in the same shaders that perform lens distortion correction and chromatic aberration correction.

Our timewarp implementation also performs speculative warping interpolated across the refresh period of the display panel. During timewarp, two distinct HMD view transformations are generated. The first transformation uses the pose at the beginning of timewarp, and the second uses the predicted pose at the end of the panel refresh period. This allows for timewarp to include IMU data even as the image is being displayed, reducing the motion-to-photon latency.

<sup>3</sup>Framebuffer objects, a method of rendering to offscreen render targets.

<sup>2</sup>Currently the open sourced version of OpenVINS does not contain loop closure, an important attribute to build reliable maps. It is therefore not a complete SLAM system with full mapping (it only does VIO). The authors plan to add loop closure in the near future. For ILLIXR, we model loop closure in the dense SLAM algorithm we use for scene reconstruction in Section IV-A2.

2) *Hologram*: Modern display optics for XR devices use rigid lenses with fixed focal lengths due to reasons of cost, weight, and size [23], [29], [54], [67]. Due to the fixed focal length, an unfortunate side-effect of such lenses is the vergence-accommodation conflict (VAC), where the user’s eyes converge at one depth, which varies with the scene, and the eye lenses focus on another, which is fixed due to the display optics [23], [29], [54], [67]. VAC is a common cause of headaches and fatigue in modern XR devices [23], [29], [54], [67].

This problem can be solved by using adaptive multi-focal displays that present several focal points instead of just one [23]. One such method of generating multiple focal planes is computational holography, wherein a phase change can be applied to each pixel using a Spatial Light Modulator (SLM) [43]<sup>4</sup>. By applying an appropriate phase change, the illusion of true depth can be created.

Ideally, we would like each pixel to be mapped to its own depth. However, this is both computationally infeasible and perceptually unnecessary. It has been shown that humans cannot perceive depths that are separated by less than 0.6 diopters [41]. Therefore, a small number of depth planes is sufficient to alleviate VAC, typically 10 to 10s of planes. One such algorithm for mapping pixels to a set of depth planes is Weighted GerchbergSaxton (GSW) [40].

GSW is an optimization algorithm that minimizes an error function based on per-pixel phase and the number of depth planes, their locations, and desired intensities. GSW performs this optimization iteratively using a forward pass from the hologram plane to the depth plane (hologram-to-depth) and a backward pass from the depth plane to the hologram plane (depth-to-hologram). The passes are repeated until a stable phase is obtained, typically in  $\sim 10$  iterations.

The forward pass propagates the phases of all the pixels to each depth plane. For a given depth plane, this includes calculating the phase contribution (a complex number) of a single pixel to that depth plane, and summing the contributions of all the pixels.

The backward pass propagates the achieved intensity of each depth plane back to each pixel. For each pixel, the intensities of the depth planes are used to obtain the new phase of the pixel, which will then be used in the next iteration of the forward pass, and so on. The fundamental computation in both passes is the evaluation of the phase equation at each pixel location [18].

We used a reference CUDA implementation of GSW as our baseline [56]. While this baseline had a high-performance implementation of GSW, it had several limitations. First, it only supported monochromatic holograms. Second, it only supported square holograms with power-of-two dimensions. Third, it was limited to holograms of size  $1024 \times 1024$  or less. We modified the CUDA kernels to support arbitrary hologram resolutions with full RGB.

### C. Audio Pipeline

Spatial audio is audio that captures the positional information of a sound source in addition to its amplitude. Spatial audio is a critical component of XR as it is imperative that the user perceive virtual sounds as real sounds, which can only be achieved via sound source localization; i.e., by using the source’s position [20].

One popular way of capturing spatial audio is Higher Order Ambisonics (HOA) [17]. HOA allows sound sources to be encoded

into a spherical soundfield which can then be manipulated and played back independently. The order of HOA determines how many virtual channels exist in the generated soundfield. For instance, order 3 HOA results in 16 virtual channels. This is an important property of HOA: the number of virtual channels in the soundfield is only dependent on the order and not on the physical number of sources that have been encoded into the field. As a result, HOA is a popular format for XR, as it can encode tens or even hundreds of sound sources without any increase in soundfield complexity [69].

Our pipeline uses libspatialaudio, an open-source HOA library [12]. We divide our pipeline into recording and playback. The two use cases are typically mutually exclusive: games have pre-encoded HOA audio and only require playback; recording in, say, a studio session only involves encoding a certain (perhaps large) number of sound sources.

1) *Recording*: We perform spatial audio recording by encoding several different monophonic sound sources simultaneously [33]. The Ambisonic encoder takes into account both the direction of the source and the attenuation and delay caused by the distance of the source. It then calculates the contribution of the source to each of the virtual sound channels in the soundfield. The final soundfield is obtained by simply adding the individual soundfields [33].

2) *Playback*: Playback is a two-step process. First, the encoded soundfield is manipulated using SLAM information. Based on the user’s head rotation and movement, the soundfield is rotated and zoomed, respectively. This step includes computing transformation matrices and applying them to each virtual channel of the soundfield [79]. Optionally, a psychoacoustic filter can be applied to the soundfield before performing the rotations in order to improve the quality of the output [17], [24], [30].

Once the desired soundfield has been obtained, binauralization maps it to the user’s headphones (HOA can decode to any number of speakers; we use headphones as they are the most common speaker configuration in XR). This involves summing the contribution of each virtual channel to the left and right speakers [28]. The most salient aspect of binauralization is the application of HRTFs or Head-Related Transfer Functions [28]. HRTFs are digital filters which capture the modification of sound by the user’s ears, head, and nose in real life. It is these modifications that allow us to localize sound so accurately in real life. The binauralization process applies HRTFs (implemented as frequency-domain convolutions) to the decoded sound in order to emulate these real-life effects [28]. There are separate HRTFs for the left and right ears due to the asymmetrical shape of the human head.

## V. EXPERIMENTAL METHODOLOGY

This section describes our experimental methodology for evaluating ILLIXR with results reported in the next section.

### A. Experimental Platform

We ran our experiments on both a desktop and an embedded class machine to understand different power performance tradeoffs for our benchmarks. Specifically, even though the desktop platform’s thermal power rating is far above what is deemed acceptable for an XR system, its performance numbers can be viewed as a rough

<sup>4</sup>A hologram is a per-pixel phase mask.






Fig. 2: Average execution time per frame, average power, and average energy per frame for all ILLIXR components on three hardware platforms.




Fig. 3: Average power breakdown for both Jetson configurations.

upper bound for CPU+GPU based near-future embedded (mobile) systems.

Both our machines consist of CPUs and GPUs. We run the CUDA and GLSL parts of our components on the GPUs and the rest on the CPU (Table II). Our desktop platform has an Intel Sandy Bridge 2600K CPU and a discrete NVIDIA Turing RTX 2060 GPU. Our embedded platform is an NVIDIA Jetson AGX Xavier development board. All experiments were run with the Jetson in 10 Watt mode, but with two different configurations; a high performance one (*Jetson-hp*) and a low power one (*Jetson-lp*), again to better understand the behavior of our benchmarks across a spectrum of power-performance tradeoffs. All clocks were maxed out in *Jetson-hp*. Clocks in *Jetson-lp* were chosen such that lowering the clocks any further would result in negligible power reduction.

### B. Execution Time Profiling

On the desktop, we used NVIDIA NSight Systems, Intel VTune, and perf to profile the CPU portion of our components. We used NSight Systems to obtain the overall execution timeline of the component, including serial CPU, parallel CPU, GPU compute, and GPU graphics phases. VTune and perf provided CPU hotspot analysis and hardware performance counter information.

To profile the GPU compute portion of our components on the desktop, we used NVIDIA NSight Compute. We were unable to profile the graphics (GLSL) portion of our components for detailed analysis as there was an unresolvable incompatibility between our hardware and NVIDIA NSight Graphics – NSight Compute’s graphics counterpart. We therefore report only end-to-end information for the graphics part.

We also instrumented each component to self-report its execution time in order to obtain execution time without profiling overhead.

For the embedded platform, we report only end-to-end execution times using the self-reported numbers of the components. We could not get more detailed breakdowns because of profiling tool limitations.

### C. Power and Energy Profiling

On the embedded platform, we collected power and energy using a custom profiler that monitored power rails to calculate both the average power and average energy for different components of the system (including CPU, GPU, DRAM, and total power).

For the desktop, we measured GPU power at a coarse level using vendor-provided tools, and CPU power using perf.

### D. Datasets

We used the following datasets for each component: Vicon Room 1 Easy, Medium, and Difficult for OpenVINS [22], dyson\_lab for ElasticFusion [5], frames from VR Museum of Fine Art [62] for the visual pipeline, and audio clips [4], [9] from Freesound [2] for the audio pipeline. None of the computations in timewarp, hologram, and audio recording and playback are input-dependent as pixel and audio sample values do not dictate what computation is performed; therefore, we found no difference between using synthetic datasets and real datasets for these components.

## VI. RESULTS

### A. Overview

Figure 2 shows the average execution time per frame, power, and energy per frame for the different ILLIXR components for the desktop, Jetson-hp, and Jetson-lp systems. (Note the log scale.) We make the following three observations.

First, the general trends follow our intuition – as we go from desktop to Jetson-hp to Jetson-lp, execution time goes up and power goes down. As expected, the desktop is significantly faster than both Jetson configurations and consumes significantly more power than both Jetson configurations. However, energy per frame presents two interesting results: the desktop is on par with Jetson for several components, and Jetson-hp consumes less energy per frame than Jetson-lp on account of being significantly more faster at a relatively lower power overhead.

Second, only some of the component-hardware configurations are able to meet their respective real-time deadlines. However, these times are only for 1) components running in isolation, and 2) in some cases (e.g., timewarp) with modern display resolutions. We anticipate that with all the components running together and with futuristic resolutions and frame rates, the already significant performance gap to real-time performance will increase further. For instance, in timewarp, a higher execution time due to more pixels would increase the motion-to-photon latency, which would degrade user experience; hologram already takes longer than its latency budget of 11 ms (at 90 Hz) on all three platforms.

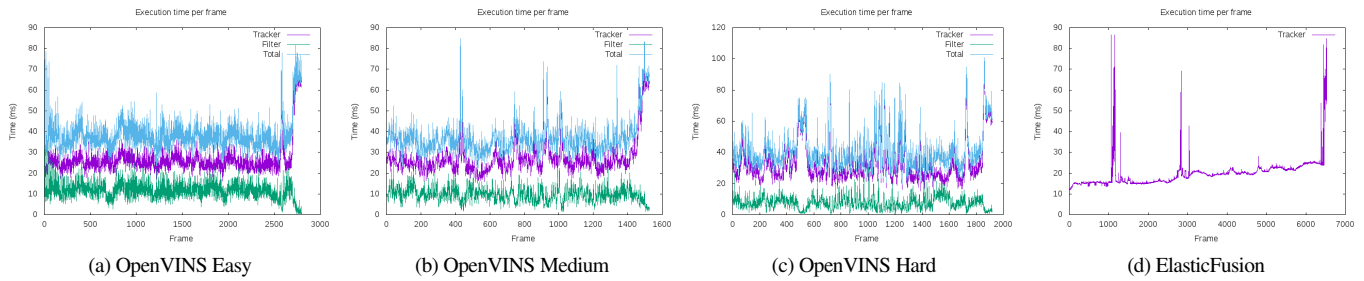


Fig. 4: Time per frame in milliseconds for both OpenVINS and ElasticFusion. For OpenVINS, the general trend is that as the scene becomes harder, tracking takes longer as features are harder to detect and match. On the other hand, the filter takes less time as there are fewer features to process. Note that in 4c almost all the time is spent in the tracker. For ElasticFusion, the spikes correspond to attempted loop closures, which take significantly more compute. Time per frame increases as time goes on due to the increasing size of the map.

Third, there is a similar gap in power consumption and energy per frame. Power is already several factors higher than the requirements stated in Table I. Energy per frame would be approximately additive when all the components are run together, and would thus increase to several *Joules* per frame. Again, larger resolutions, frame rates, and field-of-views would further widen this gap.

We believe these results clearly illustrate the need for a two to three orders of magnitude improvement in time per frame, power, and energy per frame if all day wearable XR devices are to be realized with satisfactory performance.

1) *Power Breakdown*: Figure 3 shows power for each ILLIXR component on the Jetson-hp and Jetson-lp, broken down by the hardware unit in the Jetson – CPU, GPU, DDR (DRAM), SoC (various microcontrollers and other components of the SoC), and Sys (system power including I/O). In general, the power breakdown reflects the characteristics of the individual components. CPU-only components have higher CPU power than those components that use the GPU. For components that use the GPU, GPU power is proportional to the utilization of the GPU.

Comparing Jetson-hp and Jetson-lp, the most significant change is that CPU power of all the components becomes approximately the same in lp due to a lower bound on the power consumption of Jetson. GPU power also follows a similar trend in that the difference between the various components decreases.

Figure 3 highlights the need for measuring end-to-end system power. CPU, GPU, and DDR power at best account for 77% of total power (hologram, Jetson-hp). In the worst case, they account for only 23% of the power consumption of the component (audio playback, Jetson-lp). System-level components such as on-board microcontrollers (SoC), display, disk, and other I/O blocks (Sys) are often ignored. Doing so does not capture the true power consumption of XR workloads, which typically exercise all the aforementioned components.

This can be observed by comparing OpenVINS and timewarp in Jetson-lp. The difference in SoC+Sys power between them is quite significant. The reason is the *display*. While OpenVINS is a compute-only component, timewarp is a graphics component that outputs to a display attached to the Jetson. Consequently, Sys power goes up significantly. SoC power also increases due to the on-board display controller(s).

2) *Input-Dependence*: The XR pipeline contains some components that are input-dependent and some that are not. In ILLIXR,

OpenVINS and ElasticFusion are input-dependent as the surrounding scene dictates the number of available features, difficulty of map fusion, etc. Such components present unique challenges both in terms of analysis and for specialization. To understand the input-dependent behavior, Figures 4 show per-frame execution times for OpenVINS and ElasticFusion respectively on the desktop. The figures show considerable variation in execution time both across *datasets* and across *frames*, and in fact even in *adjacent* frames. Clearly, average execution times do not faithfully capture the behavior of these components. For instance, the average for the OpenVINS medium dataset [22] is 36 ms, which is comfortably below the camera frame time of 50 ms (20 fps in our dataset). In reality, however, the frame processing time often jumps above 50 ms, up to 80 ms around frame 400. Such variation occurs in other datasets as well – in the hard dataset, the average time per frame is 39 ms, whereas most frames between frame 1100 and 1300 take significantly longer. We compare the behavior of OpenVINS with these three datasets in more detail in Section VI-C1. A similar trend can be observed for ElasticFusion, where frame processing times increase sharply during loop closure around frames 1000, 3000, and 6500.

Such behavior has several implications for architects as it makes hardware specialization more challenging. Managing thermals, performing DVFS, scheduling such computations, and sharing time and energy budgets across components are just some of the design aspects that need to be considered to efficiently deal with input-dependence. Moreover, designing hardware for the average case is not sufficient, and so mechanisms have to be developed to handle the variance in per-frame execution time.

## B. Task Diversity

Figure 5 shows the execution time breakdown of the different software tasks in each ILLIXR component on the desktop. We make three key observations.

First, there is a remarkable diversity of algorithmic tasks in each ILLIXR component, ranging from feature detection to graphics-based map fusion to signal processing and more. These tasks are differently amenable for execution on the CPU, GPU-compute, or GPU-graphics. This presents an interesting challenge for hardware specialization due to the sheer number of tasks that need to be studied.

Second, there is no single task that dominates the execution time for even a given component. The most homogeneous is audio



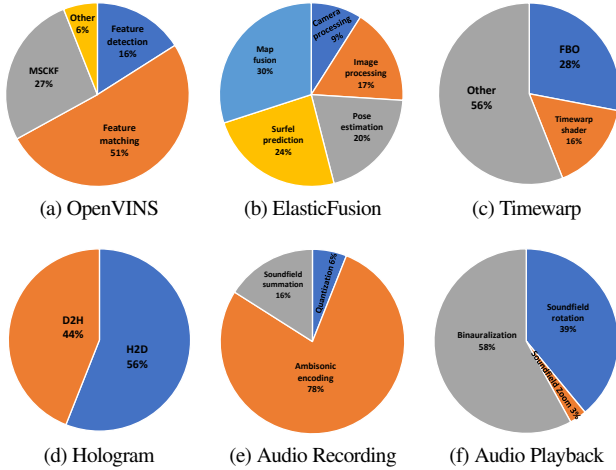


Fig. 5: Execution time breakdown of each component on the desktop. Note the diversity of sub-components and the general lack of one dominating sub-computation.

recording where ambisonic encoding is 78% of the total – but accelerating just that would limit the overall speedup to  $4.5\times$  by Amdahl’s law. ElasticFusion is the most diverse, with five major tasks and many more sub-tasks. This per-component diversity further complicates specialization as it is not sufficient to accelerate just one task even for a given component.

Third, at an algorithmic level, the tasks are typically unique to a component and not shared among components. This is yet another challenge for hardware specialization, as it is infeasible to build a unique accelerator for every task given the severe area constraints for XR devices. However, if the low-level characteristics of these tasks are similar, they may be amenable to more *general purpose* specialization. Thus, we next study these components at a more fine-grained level.

### C. Fine-grained Analysis

1) *OpenVINS*: Figure 5a shows the normalized execution time breakdown of OpenVINS with the Medium dataset, which can be first divided into the tracker, which consumes 73% of the time, and the filter, which consumes the remaining 27%. The tracker predominantly consists of feature detection (15%) and feature matching (51%), with a small percentage of time (7%) spent in miscellaneous tasks.

Both the filter and tracker are compute-intensive tasks with high locality cache accesses. This can be seen in an average IPC of 1.8 and an average L1 cache load hit rate of 97%. Feature matching only touches pixels in the neighborhood of detected features, which are typically between 200 and 300 in OpenVINS. Thus, the working set size of feature matching is small.

On the other hand, feature detection does touch every pixel in the image in order to find new features. However, there is excellent spatial locality in the task and due to the organization of the work, there is only 1 cache miss every 64 accesses (1.6% of all accesses). Furthermore, the access pattern is regular enough that it is quite likely that the prefetcher is able to eliminate the miss.

Finally, both the core kernel of feature detection, KLT, and the core kernel of MSCKF, givens rotations (GR), have slightly different characteristics than the component average. KLT and GR

have an IPC of 2.4 and 2.1, respectively, which is higher than the average IPC of 1.8. Both kernels are heavily vectorized, resulting in the higher IPC.

MSCKF, interestingly, has an L1 hit rate of 94%, non-trivially lower than the average. This is due to the fact that the size of the MSCKF feature matrix is proportional to the square of the number of features, which corresponds to 200 – 300 KB of data; i.e., larger than the 256 KB L2 of the desktop CPU. Furthermore, the memory accesses in GR are not to sequential rows – they depend on which row is being cleared out, which is input-dependent.

Across three datasets, we observe a change in the proportion of time spent in feature matching and MSCKF, as described in Section VI-A2. Comparing Easy and Hard, feature matching’s time increases from 46% to 59%, and MSCKF’s time decreases from 31% to 19%. Interestingly, we did not observe a change in average power from dataset to dataset. This is due to the fact that, as mentioned above, KLT and MSCKF have a similar IPC, and thus the difference from change in task distribution averages out.

**From an architecture perspective**, OpenVINS is challenging to specialize due to the complexity and uniqueness of each major computation. For instance, feature matching alone consists of KLT, RANSAC, and point undistortion, and each of these three sub-computations further consist of many more kernels; MSCKF consists of *seven* salient computations. On top off the complexity of designing each accelerator, the complexity of data movement between all these accelerators needs to be managed. This can be achieved by designing efficient coherence protocols that alleviate the burden of orchestrating data movement from the programmer. Such specialization can also benefit from efficient hardware design tools that increase designer productivity.

2) *ElasticFusion*: As shown in Figure 5b, ElasticFusion consists of five major tasks: processing the incoming camera feed (9%), preparing the images for pose estimation (17%), estimating the pose (20%), predicting which surfels in the map are active (24%), and fusing the new image into the existing map (30%). These tasks are implemented as a dozen CUDA kernels and almost three times as many OpenGL vertex, geometry, and fragment shaders. Therefore, we do not present an individual analysis of each but instead highlight the salient aspects of ElasticFusion.

In general, ElasticFusion is a memory-bound component. Most of the CUDA kernels use up to 200 GB/s of memory bandwidth and some surpass even 300 GB/s. This is due to several reasons.

First, there are a large number of accesses per pixel – RGB, depth, vertices, normals, cloud points, and so on. Second, different tasks have different data layouts. For instance, OpenGL stores data in RGB\_RGB format due to the nature of OpenGL vector types (which store x, y, and z sequentially). CUDA, however, prefers RR\_GG\_BB as it leads to better memory coalescing. To tackle this issue, ElasticFusion uses CUDA kernels to perform data layout transformations in between OpenGL and CUDA phases. In some instances, these kernels only access ‘RGB’ out of ‘RGBD’, which results in more data being brought into the cache than is required.

Third, cache hit rates vary significantly from kernel to kernel, and can be as low as 0% and as high as 100%. This, for instance, depends upon whether the data being accessed is laid out as RGB\_RGB or RR\_GG\_BB. In the former case, the three accesses

to R, G, and B all hit in the same cache line and thus there are a decent amount of hits. In the latter case, the hit rate is 0% as the accesses are purely streaming and no cache line is accessed twice.

Some kernels, especially image processing ones, use Sobel and Gaussian filters over windows of size  $3 \times 3$  and  $5 \times 5$ , respectively. These kernels employ 2D thread blocks, and as a result of the sliding window access pattern, achieve excellent temporal locality, which results in high cache hit rates.

In addition to its memory bound nature, ElasticFusion has two salient compute-oriented features. First, a significant amount of time is spent performing reductions. All of the pose estimation routines require global summations of some computed value, and therefore reductions are a common compute pattern. However, the reductions are over different data structures; e.g., two integers in one case and a  $6 \times 6$  matrix in another. Second, the two most expensive kernels, ICP (Iterative Closest Point) and residual calculation, have significant branch divergence due to input-dependence. In ICP, only 22 threads are active in a warp on average, which results in poor hardware utilization.

**From an architecture perspective**, ElasticFusion faces the same difficulty as OpenVINS due to the sheer number of unique sub-computations. However, the critical aspect of ElasticFusion is *data layout*. In one of the kernels, the L1 cache hit rate increased from 0% to 50% after the data layout transformation. By providing data layout transformation functionality in hardware, compute can be saved and memory accesses can be sped up. Taking this further, a more intelligent memory system could be designed via software-hardware co-design.

For instance, from one frame to another, the magnitude of movement is small (only a 33 ms gap) and so the distance between map vertices and searched vertices is small. This combined with extrapolating the next pose from a history of previous poses can be used to perform a kind of cache prefetching that not only brings in predicted data but also changes its layout in the process.

3) *Timewarp, Lens Distortion, Chromatic Aberration*: Figure 5c shows the execution time breakdown of the lens, chromatic, and timewarp component in terms of time spent setting up the FBO, running the timewarp shader, and calculating transforms and setting OpenGL state ("Other"). Out of the total 500 us, 28% are spent on the FBO setup code on the CPU and GPU, 16% running the actual timewarp shader, and 56% are spent on calculating the transformation matrices and setting OpenGL state.

The CPU code is extremely fast because the transform matrix computations are heavily vectorized, the distortion meshes are only generated *once* at the beginning of execution (and are thus not part of the main kernel), and even though there are several OpenGL calls for each frame, OpenGL command buffering keeps their overhead to a minimum. Nonetheless, a few salient OpenGL calls consume more time than others: FBO binds, clears, flushes, and texture mip-map generation. Of the two drawcalls, the first one is expensive while the second is not due to reuse of OpenGL state, parameters, and textures.

In the timewarp shader itself, all the warping calculations are performed in the vertex shader. Since there are far fewer vertices than fragments in a scene, the vertex shader is run far fewer times than the fragment shader. As a result, all the complex computations in the vertex shader only constitute a small part of the overall computation.

The fragment shader performs three texture accesses, one each for the red, green, and blue color channels. While these accesses do not all fall into the same cache line (laid out in Z-order for texture caches), the address offset is small enough in practice that they fall in nearby cache lines. As a result, texture cache locality can still be exploited. Since we did not have access to a graphics profiling tool, we confirmed this hypothesis by running the fragment shader with one, two, and three texture accesses, and measuring the difference in execution time. We did not observe an increase in execution time, although it is likely that the GPU's high texture fillrate was able to keep up with texture cache misses.

**From an architecture perspective**, timewarp has very little compute and obtains good texture cache locality once the data has been loaded, but the FBO setup and data movement itself from the framebuffer to the GPU cores is expensive. Xie et al. [75] propose an architecture that performs timewarp in memory to remove both data transfers and the overhead of managing FBOs.

4) *Hologram*: Figure 5d shows the normalized execution time breakdown of hologram in terms of its two CUDA kernels. As can be seen, both CUDA kernels are relatively well-balanced; hologram-to-depth consumes 56% of the execution time and depth-to-hologram consumes the remaining 44%. There is also a small reduction kernel in between the above two kernels but it takes less than 0.1% of the overall execution time. Both kernels are compute bound, with depth-to-hologram more so than hologram-to-depth.

**Depth-to-hologram** has an extremely high SM utilization of 90% and an extremely low memory utilization of just 5%. Almost all the executed instructions are FP fused multiply-add (FFMA), integer multiply-add (IMAD), and FP add (FADD). Consequently, the two most commonly used execution pipes are FMA and FP64. The former has a utilization of 37% while the latter has a utilization of 97%. The reason behind this is transcendental math the sine, cosine, and tangent operations that are involved in calculating the real and imaginary parts of the pixel phase. An inspection of the generated PTX showed that many of these operations were being decomposed into FP64 operations, thus the high utilization of the pipe. Furthermore, the phase calculation itself involves FP64 operations due to precision reasons.

Depth-to-hologram only uses approximately 5% of the available memory bandwidth due to an L1 cache load hit rate of 99%. Each thread has to bring in data for each depth plane to compute the new phase of the pixels the thread is responsible for. Since there are only 16 depth planes, the loads only bring in a few hundred bytes of data, and that data is reused not only by the entire thread block, but by all the other thread blocks resident on the same SM, resulting in the high L1 cache hit rate. Furthermore, the loads are all serviced from the L2. At the end of the kernel, the new phase of each pixel is written back, but that does not significantly affect performance as stores are off the critical path.

**Hologram-to-depth** has a reasonably high SM utilization of 74% and a relatively low memory utilization of 25%. Similar to depth-to-hologram, the dominant instructions are FFMA, IMAD, and FADD. However, these are not as dominant as in depth-to-hologram due to the presence of branch, sync, address calculation, and type conversion instructions. As a result, the FMA pipe has a 66% utilization, FP64 has a 74% utilization, and XU

(special function unit) has a 43% utilization. Overall SM utilization is lower than depth-to-hologram due to a tree reduction that occurs at the end of the kernel. Utilization keeps going down as the reduction progresses and active threads keep getting cut in half.

Hologram-to-depth uses significantly more memory bandwidth than depth-to-hologram 80 GB/s or 25% of the maximum available. This kernel is essentially the opposite of depth-to-hologram there are as many loads as there are pixels across all three color channels and very few stores for the depth plane data. As a result, there is significant load traffic and very little store traffic. Furthermore, each pixels phase is only used once, and thus there is no temporal locality. However, since threads in a warp access adjacent pixels, perfect coalescing is achieved and spatial locality is exploited.

**From an architecture perspective**, hologram can be accelerated via efficient transcendental hardware tightly coupled with FMA units. Depth plane data can be cheaply stored in registers and pixel data can be streamed in and out via streambuffers and FIFOs. Both kernels can easily share hardware as the nature of their computations is similar and only the flow of information is different.

5) *Audio Recording*: Figure 5e shows the normalized execution time breakdown of audio recording. Ambisonic encoding only consists of a single pipeline stage, which can be divided into three parts: conversion of 16-bit audio samples to a floating point representation, Ambisonic encoding of each monophonic sound source, and the summation of all the encodings to obtain the final soundfield.

As shown in Figure 5e, an Ambisonic encoder on average spends 6% of its execution time on sample conversion, 78% on the encoding of each sound source, and 16% on the summation of the individual soundfields. In general, all three computations are compute-bound, albeit to varying degrees. The memory footprint of encoding is small: 16 output channels each with 1024 samples only occupy 64 KB, and 32 incoming sound sources each with 1024 samples consume 128 KB. The total working set size is thus 192 KB, which can comfortably fit in the L2 cache of the CPU core. Furthermore, the loads are all streaming accesses, and thus locality need only be exploited for the stores. This can be seen in both the average load latency, which is 9 cycles, and the average memory bandwidth, which is 2.5 GB/s.

However, there are two sources of inefficiency in audio recording: one in sample conversion, and one in Ambisonic encoding. Sample conversion relies heavily on the lone, unpipelined divider unit of the CPU core. As a result, its execution is bound by instructions waiting in the queue for the divider.

Ambisonic encoding performs its multiply-accumulate operations in a column-major fashion. This is due to the fact that while memory is laid out as  $channels \times samples$ , the algorithm requires that the outer loop be over samples and the inner loop be over channels. The reason behind this is delay lines; i.e., the modelling of delay-effects in the input stream. In delay lines, the encoding of one sample affects the encoding of the next, and thus the loops cannot be interchanged. As a result, the L1 cache load hit rate is 73%, which is poor compared to traditional regular workloads. This reduces the IPC to 1.05. A possible optimization may be to lay out output memory as  $samples \times channels$ , and transpose it after the nested loops, which could be efficiently performed by a memory system that supports in-memory data layout transformations.

6) *Audio Playback*: Figure 5f shows the normalized execution time breakdown of audio playback, which can be divided into two pipeline stages: Ambisonic manipulation and binauralization. On average, the first stage consumes 42% of the total execution time (39% rotation and 3% zoom) while the second stage consumes 58% of the execution time. Audio playback has an even smaller working set size than audio recording and is even more compute bound. While the average load latency is similar to audio recording 8 cycles the average memory bandwidth is only 100 MB/s due to the smaller working set size recording had several input streams whereas playback only has a fixed Ambisonic soundfield. The access pattern is regular enough that the prefetcher can easily predict the next cache line that needs to be brought in.

The compute in audio recording is typical signal processing: frequency-domain convolutions with FIR filters implemented via FFTs and IFFTs. The huge amount of multiply-accumulate operations result in a microarchitectural utilization of 75% or higher, with the small amount of stalls being due to a few column-major accesses, use of the divider unit, and a few branch mispredictions due to the use of recursion in the FFT/IFFT implementation. The average IPC is 2.7, which is substantial, and the branch prediction accuracy is 99.7%, which reflects the regular nature of the component.

**From an architecture perspective**, Ambisonic audio processing is a compute problem with very little stress on the memory system. It can be specialized by using systolic arrays and specialized FFT blocks. Streambuffers can keep the systolic arrays and FFT blocks fully utilized given the deterministic, streaming access pattern. The output can be stored in a small scratchpad for perfect reuse.

#### D. Application-level Analysis

In order to use ILLIXR's components in a real application, the components would have to be integrated into a complete XR runtime, which is beyond the scope of this initial paper. Nevertheless, we used Snapdragon Profiler [59] to profile Minecraft Earth [50] on a Samsung Galaxy S10 [61] to perform end-to-end characterization of a modern AR application. Minecraft Earth uses ARCore [32] on Android, and thus the results reflect the behavior of ARCore's components (e.g., SLAM) and not ILLIXR's components. We were unable to obtain component-level information due to the closed-source nature of the application and runtime.

The average utilization of the CPU, GPU, and DSP in Minecraft Earth is 30%, 40%, and 12%, respectively. While the computation changes based on the scene and the user's input, the variation in utilization of the compute units is low. Since no single compute unit is fully utilized, these numbers might suggest that the system is comfortably capable of executing the application. However, the SoC temperature increases from 54 Celsius to 72 Celsius after a few minutes, which would be unacceptable in an HMD. Furthermore, modern smartphone-based AR applications are extremely simple, have poor tracking, and neither provide realistic graphics nor physics (such as occlusion).

Our experience with end-to-end profiling of modern AR applications motivates the need for an open-source XR runtime. An open-source XR runtime with state-of-the-art components would allow us to profile futuristic applications with transparency. We leave the design and implementation of such a runtime for future work.



## VII. DISCUSSION

As shown in Section VI, XR has several implications for computer architects, and specifically for specialization. First, the combined power, energy, and performance gap is **several orders of magnitude**, making this domain ideal, but challenging, for hardware specialization.

Second, reducing this gap involves addressing **system-level** components as well, such as the display and other I/O, which consume significant energy themselves. This also includes numerous sensors – while we did not model them, we expect them to be significant contributors to latency and power. For instance, shipping pixels from a camera to the SoC might consume similar energy to shipping pixels from the SoC to the display. Consequently, **on-sensor computing** would be an interesting avenue to explore.

Third, there is no single algorithmic task that dominates a particular component, and, broadly, tasks are not shared across components. While two different tasks may both be compute-intensive, the computations themselves are sufficiently different that they may need to be specialized separately. This further complicates specialization as it is in direct conflict with a core constraint of XR: limited chip area. We identify 21 tasks in Fig. 5, and expect more tasks with new components. Assuming 30 tasks and 4 accelerators per task, we get a total of 120 accelerators. Even if each accelerator occupies 2mm<sup>2</sup> and consumes 10 mW, that is an additional area of 240 mm<sup>2</sup> and additional power of 1 Watt for just compute. Interfaces between these accelerators and other peripheral logic will further add to this area and power. For reference, a modern smartphone SoC such as Apple A12 takes 83 mm<sup>2</sup> and has a TDP of 2-3 Watts [1]. Due to this constraint, **shared hardware** is almost a necessity, and therefore hardware designers have to devise ways to build such shared hardware for this diverse set of tasks.

Fourth, while we have only included one possible implementation for each of our components, there is an abundance of choices, both algorithmic and implementation, for each component. XR is a moving target and developing fixed-function hardware for a particular algorithm may not be the most cost-effective choice. For instance, for SLAM alone there are dozens of ever changing algorithms [52], [55], [57], [58], [65], [71], and designing new hardware for each is not practical. Therefore, **programmability**, to some extent, is required. This coupled with the above point leads to shared *and* programmable (yet efficient) hardware, a challenge for architects.

Finally, the input-dependence and the resulting variation in frame processing time for certain components introduces challenges of its own. These include scheduling of these components to achieve real-time deadlines under peak load, managing DVFS, and potentially sharing latency and power budgets across components. The problem is further complicated when considered in the context of shared programmable hardware – **partitioning, allocation, and scheduling of shared resources** across asynchronous tasks would be an interesting problem to tackle for computer architects and system designers alike.

We believe that ILLIXR and the characterization presented here can propel architecture research in the above directions and more; e.g., use of approximate computing, realtime scheduling environments, and more aggressive hardware-software co-design.

## VIII. RELATED WORK

There is an increasing interest in XR in the architecture community and several works have shown promising specialization results for individual components of the XR pipeline. These include Processing-in-Memory architectures for graphics [74], [75], video accelerators for VR [38], [39], [45], [46], 3D point cloud acceleration [77], stereo acceleration [27], computer vision accelerators [66], [78], and SLAM chips [44], [64]. However, these works have been for single components in isolation. Our work does not perform hardware specialization but instead provides a set of salient XR components along with insights into their characteristics. The goal of ILLIXR is to both enable XR computer architecture research and to provide an analysis of the major components of an XR pipeline.

As specialization becomes important, so does the need to better understand XR applications and workloads. There have been several works that have developed benchmark suites to address this issue. VRMark [7], FCAT [8], and Unigine [10] are examples of graphics rendering benchmarks, but do not contain the perception pipeline nor adaptive display components. The SLAMBench [19], [21], [53] series of benchmarks aid the characterization and analysis of available SLAM algorithms, but contain neither the visual pipeline nor the audio pipeline. Unlike our work, none of these benchmarks suites look at multiple components of the XR pipeline and instead just focus on one aspect of the XR pipeline.

## IX. CONCLUSION

As the need for specialization increases and architectures become increasingly domain-specific, it is important for architects to understand the requirements of emerging application domains. AR/VR/MR or XR is one such important domain. This paper presents a generic XR workflow and the first benchmark suite that represents key computations from this workflow. Our analysis shows a large number of interesting aspects for architects – demanding performance, thermal, and energy requirements; a large diversity of critical tasks such that an accelerator per task is likely to overshoot area constraints; significant input dependent computation that challenges scheduling and specialization; and a diversity in the type of instructions along with a wide variety of bottlenecks throughout the system, from graphics and compute requirements to memory bandwidth and power. ILLIXR and our analysis has the potential to propel new directions in architecture research in general, and impact the emerging domain of XR in particular. Although ILLIXR already incorporates many aspects of the XR workflow, as the industry moves towards new frontiers, such as the deep integration of machine learning and low-latency client-cloud applications, we envision ILLIXR to evolve to a standard, comprehensive, open source benchmark for XR, beyond the scope we demonstrate here. A full-stack XR runtime, open to researchers and free to users, will serve a key role in democratizing and accelerating research and analysis of both algorithms and hardware for solving XR’s most difficult challenges. ILLIXR is publicly available at <https://illixr.github.io>

## ACKNOWLEDGEMENTS

ILLIXR came together after many consultations with researchers and practitioners in many domains: audio, graphics, optics, robotics,

signal processing, and extended reality systems. We are deeply grateful for all of these discussions and specifically to the following: Wei Cu, Aleksandra Faust, Liang Gao, Matt Horsnell, Amit Jindal, Steve LaValle, Jae Lee, Steve Lovegrove, Andrew Maimone, Vegard Øye, Martin Persson, Archontis Politis, Eric Shaffer, and Paris Smaragdis.

The development of ILLIXR was supported by the Applications Driving Architectures (ADA) Research Center, a JUMP Center co-sponsored by SRC and DARPA, the Center for Future Architectures Research (C-FAR), one of the six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA, and by a Google Faculty Research Award. The development of ILLIXR was also aided by generous hardware and software donations from ARM and NVIDIA.

## REFERENCES

- [1] "A12 bionic- apple." [Online]. Available: <https://en.wikichip.org/wiki/apple/ax/a12>
- [2] "Freesound." [Online]. Available: <https://freesound.org/>
- [3] "Robot operating system." [Online]. Available: <https://www.ros.org/>
- [4] "Science Teacher Lecturing," May 2013. [Online]. Available: <https://freesound.org/people/SpliceSound/sounds/188214/>
- [5] "ElasticFusion repository," 2015. [Online]. Available: <http://github.com/mp3guy/ElasticFusion/>
- [6] "HTC VIVE Pro," 2016. [Online]. Available: <https://www.vive.com/us/product/vive-pro/>
- [7] "VRMark," 2016. [Online]. Available: <https://benchmarks.ul.com/vrmark>
- [8] "NVIDIA Frame Capture Analysis Tool," 2017. [Online]. Available: <https://www.geforce.com/hardware/technology/fgcat>
- [9] "Radio Recording," July 2017. [Online]. Available: [https://freesound.org/people/waveplay\\_old/sounds/397000/](https://freesound.org/people/waveplay_old/sounds/397000/)
- [10] "Unigine Superposition Benchmark," 2017. [Online]. Available: <https://benchmark.unigine.com/superposition>
- [11] "van Waveren Oculus Demo repository," 2017. [Online]. Available: <https://github.com/KhronosGroup/Vulkan-Samples-Deprecated/tree/master/samples/apps/atw>
- [12] "libspatialaudio repository," 2019. [Online]. Available: <https://github.com/videlolabs/libspatialaudio>
- [13] "Microsoft hololens 2," 2019. [Online]. Available: <https://www.microsoft.com/en-us/hololens/hardware>
- [14] "OpenVINS repository," 2019. [Online]. Available: [https://github.com/rpng/open\\_vins](https://github.com/rpng/open_vins)
- [15] "Snapdragon 850 mobile compute platform," Nov 2019. [Online]. Available: <https://www.qualcomm.com/products/snapdragon-850-mobile-compute-platform>
- [16] R. Aggarwal, T. P. Grantcharov, J. R. Eriksen, D. Bhirup, V. B. Kristiansen, P. Funch-Jensen, and A. Darzi, "An evidence-based virtual reality training program for novice laparoscopic surgeons," *Annals of surgery*, vol. 244, no. 2, pp. 310–314, Aug 2006, 16858196[pmid]. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/16858196>
- [17] S. Bertet, J. Daniel, E. Parizet, and O. Warusfel, "Investigation on localisation accuracy for first and higher order ambisonics reproduced sound sources," *Acta Acustica united with Acustica*, vol. 99, p. 642–657, 07 2013.
- [18] S. Bianchi and R. D. Leonardo, "Real-time optical micro-manipulation using optimized holograms generated on the gpu," *Computer Physics Communications*, vol. 181, no. 8, pp. 1444–1448, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0010465510001232>
- [19] B. Bodin, H. Wagstaff, S. Saeedi, L. Nardi, E. Vespa, J. Mawer, A. Nisbet, M. Luján, S. Furber, A. J. Davison *et al.*, "Slambench2: Multi-objective head-to-head benchmarking for visual slam," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–8.
- [20] C. H. Brown and B. J. May, *Comparative Mammalian Sound Localization*. New York, NY: Springer New York, 2005, pp. 124–178. [Online]. Available: [https://doi.org/10.1007/0-387-28863-5\\_5](https://doi.org/10.1007/0-387-28863-5_5)
- [21] M. Bujanca, P. Gafton, S. Saeedi, A. Nisbet, B. Bodin, F. O'Boyle Michael, A. J. Davison, G. Riley, B. Lennox, M. Luján *et al.*, "Slambench 3.0: Systematic automated reproducible evaluation of slam systems for robot vision challenges and scene understanding," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6351–6358.
- [22] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The euroc micro aerial vehicle datasets," *The International Journal of Robotics Research*, vol. 35, no. 10, pp. 1157–1163, 2016.
- [23] J.-H. R. Chang, B. V. K. V. Kumar, and A. C. Sankaranarayanan, "Towards multifocal displays with dense focal stacks," *ACM Trans. Graph.*, vol. 37, no. 6, pp. 198:1–198:13, Dec. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3272127.3275015>
- [24] J. Daniel, "Représentation de champs acoustiques, application la transmission et la reproduction de scènes sonores complexes dans un contexte multimedia," 01 2000.
- [25] M. S. Elbamby, C. Perfecto, M. Bennis, and K. Doppler, "Towards low-latency and ultra-reliable virtual reality," *arXiv*, 2018.
- [26] D. Evangelakos and M. Mara, "Extended timewarp latency compensation for virtual reality," in *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. ACM, 2016.
- [27] Y. Feng, P. Whatmough, and Y. Zhu, "Asv: Accelerated stereo vision system," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '52. New York, NY, USA: ACM, 2019, pp. 643–656. [Online]. Available: <http://doi.acm.org/10.1145/3352460.3358253>
- [28] M. Frank, F. Zotter, and A. Sontacchi, "Producing 3d audio in ambisonics," in *Audio Engineering Society Conference: 57th International Conference: The Future of Audio Entertainment Technology Cinema, Television and the Internet*, Mar 2015. [Online]. Available: <http://www.aes.org/e-lib/browse.cfm?elib=17605>
- [29] G. Kramida, "Resolving the vergence-accommodation conflict in head-mounted displays," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 7, pp. 1912–1931, July 2016.
- [30] M. A. Gerzon, "Practical periphony: The reproduction of full-sphere sound," in *Audio Engineering Society Convention 65*, Feb 1980. [Online]. Available: <http://www.aes.org/e-lib/browse.cfm?elib=3794>
- [31] L. Goode, "The hololens 2 puts a full-fledged computer on your face," Feb 2019. [Online]. Available: <https://www.wired.com/story/microsoft-hololens-2-headset/>
- [32] Google, "Arcore," 2019. [Online]. Available: <https://developers.google.com/ar>
- [33] F. Hollerweger, "An introduction to higher order ambisonic," 2008.
- [34] Jianbo Shi and Tomasi, "Good features to track," in *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, June 1994, pp. 593–600.
- [35] D. Kanter, "Graphics processing requirements for enabling immersive vr," *Whitepaper*, 2015.
- [36] A. Kore, "Display technologies for Augmented and Virtual Reality," 2018. [Online]. Available: <https://medium.com/inborn-experience/isplay-technologies-for-augmented-and-virtual-reality-82feca4e909f>
- [37] M. Le Chénéchal and J. Chatel-Goldman, "Htc vive pro time performance benchmark for scientific research," 2018.
- [38] Y. Leng, C.-C. Chen, Q. Sun, J. Huang, and Y. Zhu, "Semantic-aware virtual reality video streaming," in *Proceedings of the 9th Asia-Pacific Workshop on Systems*, ser. APSys '18. New York, NY, USA: ACM, 2018, pp. 21:1–21:7. [Online]. Available: <http://doi.acm.org/10.1145/3265723.3265738>
- [39] —, "Energy-efficient video processing for virtual reality," in *Proceedings of the 46th International Symposium on Computer Architecture*, ser. ISCA '19. New York, NY, USA: ACM, 2019, pp. 91–103. [Online]. Available: <http://doi.acm.org/10.1145/3307650.3322264>
- [40] R. D. Leonardo, F. Ianni, and G. Ruocco, "Computer generation of optimal holograms for optical trap arrays," *Opt. Express*, vol. 15, no. 4, pp. 1913–1922, Feb 2007. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-15-4-1913>
- [41] S. Liu and H. Hua, "A systematic method for designing depth-fused multi-focal plane three-dimensional displays," *Opt. Express*, vol. 18, no. 11, pp. 11562–11573, May 2010. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-18-11-11562>
- [42] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2*, ser. IJCAI'81. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1981, pp. 674–679. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1623264.1623280>
- [43] A. Maimone, A. Georgiou, and J. S. Kollin, "Holographic near-eye displays for virtual and augmented reality," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, p. 85, 2017.
- [44] D. K. Mandal, S. Jandhyala, O. J. Omer, G. S. Kalsi, B. George, G. Neela, S. K. Rethinagiri, S. Subramoney, L. Hacking, J. Radford, E. Jones, B. Kuttanna, and H. Wang, "Visual inertial odometry at the edge: A hardware-software co-design approach for ultra-low latency and power," in *2019 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2019, pp. 960–963.

- [45] A. Mazumdar, T. Moreau, S. Kim, M. Cowan, A. Alaghi, L. Ceze, M. Oskin, and V. Sathe, "Exploring computation-communication tradeoffs in camera systems," in *2017 IEEE International Symposium on Workload Characterization (IISWC)*, Oct 2017, pp. 177–186.
- [46] A. Mazumdar, A. Alaghi, J. T. Barron, D. Gallup, L. Ceze, M. Oskin, and S. M. Seitz, "A hardware-friendly bilateral solver for real-time virtual reality video," in *Proceedings of High Performance Graphics*, ser. HPG '17. New York, NY, USA: ACM, 2017, pp. 13:1–13:10. [Online]. Available: <http://doi.acm.org/10.1145/3105762.3105772>
- [47] M. McGuire, "Exclusive: How nvidia research is reinventing the display pipeline for the future of vr, part 1," 2017. [Online]. Available: <https://www.roadtovr.com/exclusive-how-nvidia-research-is-reinventing-the-display-pipeline-for-the-future-of-vr-part-1/>
- [48] —, "Exclusive: How nvidia research is reinventing the display pipeline for the future of vr, part 2," 2017. [Online]. Available: <https://www.roadtovr.com/exclusive-nvidia-research-reinventing-display-pipeline-future-vr-part-2/>
- [49] L. McMillan and G. Bishop, "Head-tracked stereoscopic display using image warping," in *Stereoscopic Displays and Virtual Reality Systems II*, S. S. Fisher, J. O. Merritt, and M. T. Bolas, Eds., vol. 2409, International Society for Optics and Photonics. SPIE, 1995, pp. 21 – 30. [Online]. Available: <https://doi.org/10.1117/12.205865>
- [50] Mojang, "Minecraft earth," 2019. [Online]. Available: <https://www.minecraft.net/en-us/about-earth/>
- [51] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint kalman filter for vision-aided inertial navigation," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, April 2007, pp. 3565–3572.
- [52] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: a versatile and accurate monocular SLAM system," *CoRR*, vol. abs/1502.00956, 2015. [Online]. Available: <http://arxiv.org/abs/1502.00956>
- [53] L. Nardi, B. Bodin, M. Z. Zia, J. Mawer, A. Nisbet, P. H. Kelly, A. J. Davison, M. Luján, M. F. O'Boyle, G. Riley *et al.*, "Introducing slambench, a performance and accuracy benchmarking methodology for slam," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 5783–5790.
- [54] N. Padmanaban, R. Konrad, E. A. Cooper, and G. Wetzstein, "Optimizing vr for all users through adaptive focus displays," in *ACM SIGGRAPH 2017 Talks*, ser. SIGGRAPH '17. New York, NY, USA: ACM, 2017, pp. 77:1–77:2. [Online]. Available: <http://doi.acm.org/10.1145/3084363.3085029>
- [55] W. L. Y. Y. Patrick Geneva, Kevin Eckenhoff and G. Huang, "Openvins: A research platform for visual-inertial estimation," *IROS 2019 Workshop on Visual-Inertial Navigation: Challenges and Applications*, 2019.
- [56] M. Persson, D. Engström, and M. Goksr, "Real-time generation of fully optimized holograms for optical trapping applications," in *Optical Trapping and Optical Micromanipulation VIII*, K. Dholakia and G. C. Spalding, Eds., vol. 8097, International Society for Optics and Photonics. SPIE, 2011, pp. 291 – 299. [Online]. Available: <https://doi.org/10.1117/12.893599>
- [57] T. Qin, S. Cao, J. Pan, and S. Shen, "A general optimization-based framework for global pose estimation with multiple sensors," *arXiv preprint arXiv:1901.03642*, 2019.
- [58] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [59] Qualcomm, "Snapdragon profiler," 2019. [Online]. Available: <https://developer.qualcomm.com/software/snapdragon-profiler>
- [60] E. Rosten and T. Drummond, "Machine learning for high-speed corner detection," in *Proceedings of the 9th European Conference on Computer Vision - Volume Part I*, ser. ECCV'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 430–443. [Online]. Available: [http://dx.doi.org/10.1007/11744023\\_34](http://dx.doi.org/10.1007/11744023_34)
- [61] Samsung, "Samsung galaxy s10," 2019. [Online]. Available: <https://www.samsung.com/us/mobile/galaxy-s10/>
- [62] F. Sinclair, "The vr museum of fine art," 2016. [Online]. Available: [https://store.steampowered.com/app/515020/The\\_VR\\_Museum\\_of\\_Fine\\_Art/](https://store.steampowered.com/app/515020/The_VR_Museum_of_Fine_Art/)
- [63] Skarredghost, "All you need to know on hololens 2," May 2019. [Online]. Available: <https://skarredghost.com/2019/02/24/all-you-need-know-hololens-2/>
- [64] A. Suleiman, Z. Zhang, L. Carlone, S. Karaman, and V. Sze, "Navion: A 2-mw fully integrated real-time visual-inertial odometry accelerator for autonomous navigation of nano drones," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 4, pp. 1106–1119, April 2019.
- [65] K. Sun, K. Mohta, B. Pfrommer, M. Watterson, S. Liu, Y. Mulgaonkar, C. J. Taylor, and V. Kumar, "Robust stereo visual inertial odometry for fast autonomous flight," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 965–972, 2018.
- [66] S. Triest, D. Nikolov, J. Rolland, and Y. Zhu, "Co-optimization of optics, architecture, and vision algorithms," in *Workshop on Approximate Computing Across the Stack (WAX)*, 2019.
- [67] P. R. K. Turnbull and J. R. Phillips, "Ocular effects of virtual reality headset wear in young adults," in *Scientific Reports*, 2017.
- [68] J. van Waveren, "The asynchronous time warp for virtual reality on consumer hardware," in *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*. ACM, 2016, pp. 37–46.
- [69] A. Vilkaitis, M. Dring, A. Hill, C. Middlicott, and B. Wiggins, "Room acoustics and virtual reality: An implementation of auralisation and 360 degree image techniques to create virtual representations of spaces," 11 2016.
- [70] D. Wagner, "Motion to Photon Latency in Mobile AR and VR," 2018. [Online]. Available: <https://medium.com/@DAQRI/motion-to-photon-latency-in-mobile-ar-and-vr-99f82c480926>
- [71] T. Whelan, S. Leutenegger, R. Salas-Moreno, B. Glocker, and A. Davison, "Elasticfusion: Dense slam without a pose graph." *Robotics: Science and Systems*, 2015.
- [72] Wikipedia contributors, "Hololens 2 — Wikipedia, the free encyclopedia," 2019. [Online]. Available: [https://en.wikipedia.org/wiki/HoloLens\\_2](https://en.wikipedia.org/wiki/HoloLens_2)
- [73] —, "Htc vive — Wikipedia, the free encyclopedia," 2019. [Online]. Available: [https://en.wikipedia.org/w/index.php?title=HTC\\_Vive&oldid=927924955](https://en.wikipedia.org/w/index.php?title=HTC_Vive&oldid=927924955)
- [74] C. Xie, S. L. Song, J. Wang, W. Zhang, and X. Fu, "Processing-in-memory enabled graphics processors for 3d rendering," in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Feb 2017, pp. 637–648.
- [75] C. Xie, X. Zhang, A. Li, X. Fu, and S. Song, "Pim-vr: Erasing motion anomalies in highly-interactive virtual reality world with customized memory cube," in *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Feb 2019, pp. 609–622.
- [76] XinReality, "Microsoft hololens," 2017. [Online]. Available: [https://xinreality.com/wiki/Microsoft\\_HoloLens](https://xinreality.com/wiki/Microsoft_HoloLens)
- [77] T. Xu, B. Tian, and Y. Zhu, "Tigris: Architecture and algorithms for 3d perception in point clouds," in *Proceedings of the 52Nd Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '52. New York, NY, USA: ACM, 2019, pp. 629–642. [Online]. Available: <http://doi.acm.org/10.1145/3352460.3358259>
- [78] Y. Zhu, A. Samajdar, M. Mattina, and P. Whatmough, "Euphrates: Algorithm-soc co-design for low-power mobile continuous vision," in *Proceedings of the 45th Annual International Symposium on Computer Architecture*, ser. ISCA '18. Piscataway, NJ, USA: IEEE Press, 2018, pp. 547–560. [Online]. Available: <https://doi.org/10.1109/ISCA.2018.00052>
- [79] F. Zotter and M. Frank, "All-round ambisonic panning and decoding," *J. Audio Eng. Soc.*, vol. 60, no. 10, pp. 807–820, 2012. [Online]. Available: <http://www.aes.org/e-lib/browse.cfm?elib=16554>