

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

Muhammad Huzaifa Rishi Desai Samuel Grayson Xutao Jiang Ying Jing
Jae Lee Fang Lu Yihan Pang Joseph Ravichandran Finn Sinclair
Boyuan Tian Hengzhi Yuan Jeffrey Zhang Sarita V. Adve
University of Illinois at Urbana-Champaign
illixr@cs.illinois.edu

ABSTRACT

As we enter the era of domain-specific architectures, systems researchers must 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 ILLIXR, the first open source end-to-end XR system (1) with state-of-the-art components, (2) integrated with a modular and extensible multithreaded runtime, (3) providing an OpenXR compliant interface to XR applications (e.g., game engines), and (4) with the ability to report (and trade off) several quality of experience (QoE) metrics. We analyze performance, power, and QoE metrics for the complete ILLIXR system and for its individual components. Our analysis reveals several properties with implications for architecture and systems research. These include demanding performance, power, and QoE requirements, a large diversity of critical tasks, inter-dependent execution pipelines with challenges in scheduling and resource management, and a large tradeoff space between performance/power and human perception related QoE metrics. ILLIXR and our analysis have the potential to propel new directions in architecture and systems research in general, and impact XR in particular.

1. INTRODUCTION

Recent years have seen the convergence of multiple disruptive trends to fundamentally change computer systems: (1) With the end of Dennard scaling and Moore’s law, application-driven 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 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 to collectively meet end-user demands.

Case for Extended Reality (XR) as a driving domain: This

paper makes the case that the emerging domain of virtual, augmented, and mixed reality, collectively referred to as extended reality (XR),¹ is a rich domain that can propel research on efficient domain-specific edge systems. Our case rests on the following observations:

(1) *Pervasive:* XR will pervade most aspects of our lives — it will affect the way we teach, conduct science, practice medicine, entertain ourselves, train professionals, interact socially, and more. Indeed, XR is envisioned to be the next interface for most of computing [2, 54, 72, 115].

(2) *Challenging demands:* While current XR systems exist today, they are far from providing a tetherless experience approaching perceptual abilities of humans. There is a gap of several orders of magnitude between what is needed and achievable in performance, power, and usability (Table 1), giving architects a potentially rich space to innovate.

(3) *Multiple and diverse components:* XR involves a number of diverse sub-domains — video, graphics, computer vision, machine learning, optics, audio, and components of robotics — making it challenging to design a system that executes each one well while respecting the resource constraints.

(4) *Full-stack implications:* The combination of real-time constraints, complex interacting pipelines, and ever-changing algorithms creates a need for full stack optimizations involving the hardware, compiler, operating system, and algorithm.

(5) *Flexible accuracy for end-to-end user experience:* The end user being a human with limited perception enables a rich space of accuracy-aware resource trade-offs, but requires the ability to quantify impact on end-to-end experience.

Case for an XR system testbed: A key obstacle to architecture research for XR is that there are no open source benchmarks covering the entire XR workflow to drive such research. While there exist open source codes for some individual components of the XR workflow (typically developed by domain researchers), there is no integrated suite that enables researching an XR system. As we move from the era of general-purpose, homogeneous cores on chip to domain-specific, heterogeneous *system-on-chip* architectures, benchmarks need to follow the same trajectory. While previous benchmarks comprising of suites of independent applications (e.g., Parsec [7], Rodinia [16], SPEC [10, 39], SPLASH [100, 103, 124], and more) sufficed to evaluate general-purpose single- and multicore architectures, there

¹Virtual reality (VR) immerses the user in a completely digital environment. Augmented Reality (AR) enhances the user’s real world with overlaid digital content. Mixed reality (MR) goes beyond AR in enabling the user to interact with virtual objects in their real world.

is now a need for a *full-system-benchmark* methodology, better viewed as a *full system testbed*, to design and evaluate system-on-chip architectures. Such a methodology must bring together the diversity of components that will interact with each other in the domain-specific system and also be flexible and extensible to accept future new components.

An XR full-system-benchmark or testbed will continue to enable traditional research of designing optimized accelerators for a given XR component with conventional metrics such as power, performance, and area for that component, but additionally allowing evaluations for the end-to-end impact on the system. More important, the integrated collection of components in a full XR system will enable new research that co-designs acceleration for the multiple diverse and demanding components of an XR system, across the full stack, and with end-to-end user experience as the metric.

Challenges and contributions: This paper presents ILLIXR, the first open-source XR full-system testbed; an analysis of performance, power, and QoE metrics for ILLIXR on a desktop class and embedded class machine; and several implications for future systems research.

There were two challenges in the development of ILLIXR. First, ILLIXR required expertise in a large number of sub-domains (e.g., robotics, computer vision, graphics, optics, and audio). We developed ILLIXR after consultation with many experts in these sub-domains, which led to identifying a representative XR workflow and state-of-the-art algorithms and open source codes for the constituent components.

Second, until recently, commercial XR devices had proprietary interfaces. For example, the interface between an Oculus head mounted device (HMD) runtime and the Unity or Unreal game engines that run on the HMD has been closed as are the interfaces between the different components within the HMD runtime. OpenXR [111], an open standard, was released in July 2019 to partly address this problem by providing a standard for XR device runtimes to interface with the applications that run on them. ILLIXR leverages an open implementation of this new standard [77] to provide an OpenXR compliant XR system. However, OpenXR is still evolving, it does not address the problem of interfaces between the different components within an XR system, and the ecosystem of OpenXR compliant applications (e.g., game engines and games) is still developing. Nevertheless ILLIXR is able to provide the functionality to support sophisticated applications.

Specifically, this work makes the following contributions.

(1) We develop ILLIXR, the first open source XR system, consisting of a representative XR workflow (i) with state-of-the-art components, (ii) integrated with a modular and extensible multithreaded runtime, (iii) providing an OpenXR compliant interface to XR applications (e.g., game engines), and (iv) with the ability to report (and trade off) several quality of experience (QoE) metrics. The ILLIXR components represent the sub-domains of robotics (odometry or SLAM), computer vision (scene reconstruction), machine learning (eye tracking), image processing (camera), graphics (asynchronous reprojection), optics (lens distortion and chromatic aberration correction), audio (3D audio encoding and decoding), and displays (holographic displays).

(2) We analyze performance, power, and QoE metrics for

ILLIXR on desktop and embedded class machines with CPUs and GPUs, driven by a game engine running representative VR and AR applications. Overall, we find that current systems are far from the needs of future devices, making the case for efficiency through techniques such as specialization, codesign, and approximation.

(3) Our system level analysis shows a wide variation in the resource utilization of different components. Although components with high resource utilization should clearly be targeted for optimization, components with relatively low utilization are also critical due to their impact on QoE.

(4) Power breakdowns show that CPUs, GPUs, and memories are only part of the power consumption. The rest of the SoC and system logic, including data movement for displays and sensors is a major component as well, motivating technologies such as on-sensor computing.

(5) We find XR components are quite diverse in their use of CPU, GPU compute, and GPU graphics, and exhibit a range of IPC and system bottlenecks. Analyzing their compute and memory characteristics, we find a variety of patterns and subtasks, with none dominating. The number and diversity of these patterns poses a research question for the granularity at which accelerators should be designed and whether and how they should be shared among different components. These observations motivate research in automated tools to identify acceleratable primitives, architectures for communication between accelerators, and accelerator software interfaces and programming models.

(6) Most components exhibit significant variability in per-frame execution time due to input-dependence or resource contention, thereby making it challenging to schedule and allocate shared resources. Further, there are a large number of system parameters that need to be tuned for an optimal XR experience. The current process is mostly ad hoc and exacerbated by the above variability.

Overall, this work provides the architecture and systems community with a one-of-a-kind infrastructure, foundational analyses, and case studies to enable new research within each layer of the system stack or co-designed across the stack, impacting a variety of domain-specific design methodologies in general and for XR in particular. ILLIXR is fully open source and is designed to be extensible so that the broader community can easily contribute new components and features and new implementations of current ones. The current ILLIXR codebase consists of over 100K lines of code for the main components, another 80K for the OpenXR interface implementation, several XR applications and an open source game engine, and continues to grow with new components and algorithms. ILLIXR can be found at:

<https://illixr.github.io>

2. XR REQUIREMENTS AND METRICS

XR requires devices that are lightweight, mobile, and all day wearable, and provide sufficient compute at low thermal and energy consumption to support a high quality of experience (QoE). Table 1 summarizes values for various system-level quality related metrics for current state-of-the-art XR devices and the aspiration for ideal futuristic devices. These requirements are driven by both human anatomy and usage considerations; e.g., ultra-high performance to emulate

Table 1: Ideal requirements of VR and AR vs. state-of-the-art devices, HTC Vive Pro for VR and Microsoft HoloLens 2 for AR. VR devices are typically larger and so afford more power and thermal headroom. The HTC Vive Pro offloads most of its work to an attached server, so its power and area values are not meaningful. The ideal case requirements for power, area, and weight are based on current devices which are considered close to ideal in these (but not other) respects – Snapdragon 835 in the Oculus Quest VR headset and APQ8009w in the North Focals small AR glasses.

Metric	HTC Vive Pro	Ideal VR [22, 49]	Microsoft HoloLens 2	Ideal AR [22, 49, 110]
Resolution (MPixels)	4.6 [44]	200	4.4 [76]	200
Field-of-view (Degrees)	110 [44]	Full:	52 diagonal	Full:
		165×175	[36, 122]	165×175
		Stereo:		Stereo:
		120×135		120×135
Refresh rate (Hz)	90 [44]	90 – 144	120 [104]	90 – 144
Motion-to-photon latency (ms)	< 20 [18]	< 20	< 9 [97]	< 5
Power (W)	N/A	1 – 2	> 7 [41, 46, 105]	0.1 – 0.2
Silicon area (mm ²)	N/A	100 – 200	> 173 [19, 46]	< 100
Weight (grams)	470 [123]	100 – 200	566 [76]	10s

the real world, extreme thermal constraints due to contact with human skin, and light weight for comfort. We identified the values of the various aspirational metrics through an extensive survey of the literature [49, 72, 73]. 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 AR 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.

3. THE ILLIXR SYSTEM

Figure 1 presents the ILLIXR system. ILLIXR captures components that belong to an XR runtime such as Oculus VR (OVR) and are shipped with an XR headset. Applications, often built using a game engine, are separate from the XR system or runtime, and interface with it using the runtime’s API. For example, a game developed on the Unity game engine for an Oculus headset runs on top of OVR, querying information from OVR and submitting rendered frames to it using the OVR API. Analogously, applications interface with ILLIXR using the OpenXR API [111] (we leverage the OpenXR API implementation from MonoDro [77]). Consequently, ILLIXR does not include components from the application, but captures the performance impact of the application running on ILLIXR. We collectively refer to all application-level tasks such as user input handling, scene simulation, physics, rendering, etc. as "application" in the rest of this paper.

As shown in Figure 1, ILLIXR’s workflow consists of three pipelines – perception, visual, and audio – each containing several components and interacting with each other through the ILLIXR communication interface and runtime. Figure 2 illustrates these interactions – the left side presents a timeline for an ideal schedule for the different components and the right side shows the dependencies between the different components (enforced by the ILLIXR runtime). We distilled this workflow from multiple sources that describe different parts of a typical VR, AR, or MR pipeline, including conversations with several experts from academia and industry.

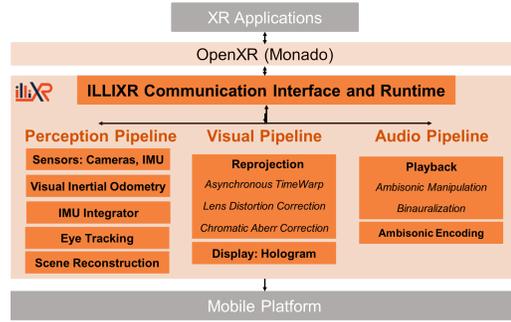


Figure 1: The ILLIXR system and its relation to XR applications, the OpenXR interface, and mobile platforms.

ILLIXR represents a state-of-the-art system capable of running typical XR applications and providing an end-to-end XR user experience. Nevertheless, XR is an emerging and evolving domain and no XR experimental testbed or commercial device can be construed as complete in the traditional sense. Compared to specific commercial headsets today, ILLIXR may miss a component (e.g., Oculus Quest 2 provides hand tracking) and/or it may have additional or more advanced components (e.g., except for HoloLens 2, most current XR systems do not have scene reconstruction). Further, while ILLIXR supports state-of-the-art algorithms for its components, new algorithms are continuously evolving. ILLIXR therefore supports a modular and extensible design that makes it relatively easy to swap and to add new components.

The system presented here assumes all computation is done at the device (i.e., no offloading to the cloud, server, or other edge devices) and we assume a single user (i.e., no communication with multiple XR devices). Thus, our workflow does not represent any network communication. We are currently working on integrating support for networking, edge and cloud work partitioning, and multiparty XR within ILLIXR, but these are outside the scope of this paper.

Sections 3.1– 3.3 next describe the three ILLIXR pipelines and their components, Section 3.4 describes the modular runtime architecture that integrates these components, and Section 3.5 describes the metrics and telemetry support in ILLIXR. Table 2 summarizes the algorithm and implementation information for each component in the three pipelines. ILLIXR already supports multiple, easily interchangeable alternatives for some components; for lack of space, we pick one alternative, indicated by a * in the table, for detailed results in this paper. Tables 3–8 summarize, for each for these components, their major tasks and sub-tasks in terms of their functionality, computation and memory patterns, and percentage contribution to execution time, discussed further below and in Section 5 (we do not show the sensor related ILLIXR components and IMU integrator since they are quite simple).

3.1 Perception Pipeline

The perception pipeline translates the user’s physical motion 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. The input to this part comes from sensors such as cameras and an inertial measurement unit (IMU)

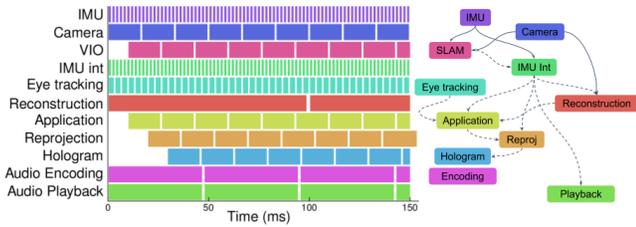


Figure 2: Interactions between ILLIXR components. The left part shows an ideal ILLIXR execution schedule and the right part shows inter-component dependencies that the ILLIXR scheduler must maintain (Section 3.4). Solid arrows are synchronous and dashed are asynchronous dependencies.

Table 2: ILLIXR component algorithms and implementations. GLSL stands for OpenGL Shading Language. * represents the implementation alternative for which we provide detailed results.

Component	Algorithm	Implementation
Perception Pipeline		
Camera	ZED SDK* [107]	C++
Camera	Intel RealSense SDK [47]	C++
IMU	ZED SDK [107]	C++
IMU	Intel RealSense SDK [47]	C++
VIO	OpenVINS* [92]	C++
VIO	Kimera-VIO [51]	C++
IMU Integrator	RK4* [92]	C++
IMU Integrator	GTSAM [37]	C++
Eye Tracking	RITnet [15]	Python,CUDA
Scene Reconstruction	ElasticFusion* [21]	C++,CUDA,GLSL
Scene Reconstruction	KinectFusion [52]	C++, CUDA
Visual Pipeline		
Reprojection	VP-matrix reprojection w/ pose [116]	C++, GLSL
Lens Distortion	Mesh-based radial distortion [116]	C++, GLSL
Chromatic Aberration	Mesh-based radial distortion [116]	C++, GLSL
Adaptive display	Weighted Gerchberg–Saxton [94]	CUDA
Audio Pipeline		
Audio Encoding	Ambisonic encoding [62]	C++
Audio Playback	Ambisonic manipulation, binauralization [62]	C++

to provide the user’s acceleration and angular velocity. The processing components include camera and IMU processing, head tracking or VIO (Visual Inertial Odometry) for obtaining low frequency but precise estimates of the user’s pose (the position and orientation of their head), IMU integration for obtaining high frequency pose estimates, eye tracking for determining the gaze of the user, and scene reconstruction for generating a 3D model of the user’s surroundings. We elaborate on some of the key components below.

3.1.1 Camera

ILLIXR can work with any commercial camera. We provide support for both ZED Mini and Intel RealSense cameras. For this paper, we use a ZED Mini [106] as it provides IMU readings, stereo RGB images, and depth images all together with a simple API. The camera component first asks the camera to obtain new stereo camera images and then copies these images into main memory.

3.1.2 IMU

ILLIXR’s IMU component queries the IMU at fixed intervals to obtain linear acceleration from the accelerometer and angular velocity from the gyroscope. In addition to retrieving the IMU values, this component converts the angular velocity

from deg/sec to rad/sec , and timestamps the readings for use by other components, such as VIO and IMU integrator.

3.1.3 VIO

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, and provide both the position and orientation of the user’s head; i.e., the six degrees of freedom or 6DOF pose of the user. There are a large number of algorithms that have been proposed for SLAM for a variety of scenarios, including for robots, drones, VR, and AR. We support both OpenVINS [33] and Kimera-VIO [99]. In this paper, we use OpenVINS as it has been shown to be more accurate than other popular algorithms such as VINS-Mono, VINS-Fusion, and OKVIS [33]. Both OpenVINS and Kimera-VIO do not have mapping capabilities, and are therefore called Visual-Inertial Odometry (VIO) algorithms. Table 3 overviews OpenVINS by describing its tasks.²

3.1.4 IMU Integrator

VIO only generates poses at the rate of the camera, which is typically low. To obtain high speed estimates of the user’s pose, an IMU integrator integrates incoming IMU samples from the IMU to obtain a relative pose and adds it to the VIO’s latest pose to obtain the current pose of the user. ILLIXR has two IMU integrators available: one utilizes GTSAM to perform integration [13, 26] and the other uses RK4, inspired by OpenVINS’ IMU integrator [33]. We use the RK4 integrator in this paper.

3.1.5 Eye Tracking

We used RITnet for eye-tracking [15] as it won the OpenEDS eye tracking challenge [32] with an accuracy of 95.3%. RITnet is a deep neural network that combines U-Net [98] and DenseNet [45]. The network takes as input a grayscale eye image and segments it into the background, iris, sclera, and pupil. Since neural networks are well understood in the architecture community, we omit the task table for space.

3.1.6 Scene Reconstruction

For scene reconstruction, we chose KinectFusion [83] and ElasticFusion [121], two dense SLAM algorithms that use an RGB-Depth (RGB-D) camera to build a dense 3D map of the world. In this paper, we focus on ElasticFusion (Table 4), which uses a surfel-based map representation, where each map point, called a surfel, consists of a position, normal, color, radius, and timestamp. In each frame, incoming color and depth data from the camera is used to update the relevant surfels in the map.

3.2 Visual Pipeline

The visual pipeline takes information about the user’s new pose from the perception pipeline, the submitted frame from the application, and produces the final display using two parts: asynchronous reprojection and display.

3.2.1 Asynchronous Reprojection

²Due to space constraints, we are unable to describe here the details of the algorithms in the components of ILLIXR.

Table 3: Task breakdown of VIO.

Task	Time	Computation	Memory Pattern
Feature detection Detects new features in the new camera images	15%	Integer stencils per pyramid level	Locally dense stencil; globally mixed dense and sparse
Feature matching Matches features across images	13%	Integer stencils; GEMM; linear algebra	Locally dense stencil; globally mixed dense and sparse; mixed dense and random feature map accesses
Filter Estimates 6DOF pose using camera and IMU measurements	62%	Gauss-Newton refinement; QR decomposition; GEMM; linear algebra	Mixed dense and sparse feature map and filter matrix accesses
Other Miscellaneous tasks	10%	Gaussian filter; histogram	Globally dense stencil

Table 4: Task breakdown of scene reconstruction.

Task	Time	Computation	Memory Pattern
Camera Processing Processes incoming camera depth image	5%	Bilateral filter; invalid depth rejection	Dense sequential accesses to depth image
Image Processing Pre-processes RGB-D image for tracking and mapping	18%	Generation of vertex map, normal map, and image intensity; image undistortion; pose transformation of old map	Globally dense; local stencil; layout change from RGB_RGB → RR_GG_BB;
Pose Estimation Estimates 6DOF pose	28%	ICP; photometric error; geometric error	photometric error is globally dense; others are globally sparse, locally dense
Surfel Prediction Calculates active surfels in current frame	38%	Vertex and fragment shaders	Globally sparse; locally dense
Map Fusion Updates map with new surfel information	11%	Vertex and fragment shaders	Globally sparse; locally dense

Table 5: Task breakdown of reprojection.

Task	Time	Computation	Memory Pattern
FBO FBO state management	24%	Framebuffer bind and clear	Driver calls; CPU-GPU communication
OpenGL State Update Sets up OpenGL state	54%	OpenGL state updates; one drawcall per eye	Driver calls; CPU-GPU communication
Reprojection Applies reprojection transformation to image	22%	6 matrix-vector MULs/vertex	Accesses uniform, vertex, and fragment buffers; 3 texture accesses/fragment

Asynchronous reprojection corrects the rendered image submitted by the application for optical distortions and compensates for latency from the rendering process (which adds a lag between the user’s movements and the delivery of the image to the user). The reprojection component reprojects the latest available frame based on a prediction of the user’s

Table 6: Task breakdown of hologram.

Task	Time	Computation	Memory Pattern
Hologram-to-depth Propagates pixel phase to depth plane	57%	Transcendentals; FMADDS; TB-wide tree reduction	Dense row-major; spatial locality in pixel data; temporal locality in depth data; reduction in scratchpad
Sum Sums phase differences from hologram-to-depth	< 0.1%	Tree reduction	Dense row-major; reduction in scratchpad
Depth-to-hologram Propagates depth plane phase to pixel	43%	Transcendentals; FMADDS; thread-local reduction	Dense row-major; no pixel reads; pixels written once

movement, resulting in less perceived latency. The prediction is performed by obtaining the latest pose from the IMU integrator and extrapolating it to the display time of the image using a constant acceleration model. Reprojection is also used to compensate for frames that might otherwise be dropped because the application could not submit a frame on time to the runtime. Reprojection is critical in XR applications, as any perceived latency will cause discomfort [23, 74, 118].

In addition to latency compensation, the reprojection component performs lens distortion correction and chromatic aberration correction. Most XR 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 [53]. These optics usually induce significant optical distortion and chromatic aberration – lines become curved and colors are non-uniformly refracted. These optical affects are corrected by applying reverse distortions to the rendered image, such that the distortion from the optics is mitigated [116].

Our reprojection, lens distortion correction, and chromatic aberration correction implementation is derived from [120], which only supports rotational reprojection.³ We extracted the distortion and aberration correction, and reprojection shaders from this reference, and implemented our own version in OpenGL that integrates reprojection, lens distortion correction, and chromatic aberration correction into one component (Table 5).

3.2.2 Display

The fixed focal length of modern display optics causes a 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. VAC is a common cause of headaches and fatigue in modern XR devices [14, 31, 93, 113]. One possible solution to this problem is computational holography, wherein a phase change is applied to each pixel using a Spatial Light Modulator [68] to generate several focal points instead of just one. Since humans cannot perceive depths less than 0.6 diopters apart [65], typically 10 or so depth planes are sufficient. The per-pixel phase mask is called a hologram.

In ILLIXR, we use Weighted Gerchberg–Saxton (GSW) [60] to generate holograms (Table 6). We used a reference CUDA implementation of GSW [94], and extended it to support RGB holograms of arbitrary size.

³We are currently developing positional reprojection.

Table 7: Task breakdown of audio encoding.

Task	Time	Computation	Memory Pattern
Normalization INT16 to FP32	7%	Element-wise FP32 division	Dense row-major
Encoding Sample to soundfield mapping	81%	$Y[j][i] = D \times X[j]$	Dense column-major
Summation HOA soundfield summation	11%	$Y[i][j] = \sum_k X_k[i][j] \forall k$	Dense row-major

3.3 Audio Pipeline

The audio pipeline is responsible for generating realistic spatial audio and is composed of audio encoding and playback (Table 7 and Table 8). Our audio pipeline uses libspatialaudio, an open-source Higher Order Ambisonics library [62].

3.3.1 Encoding

We perform spatial audio encoding by encoding several different monophonic sound sources simultaneously into an Ambisonic soundfield [42].

3.3.2 Playback

During playback, Ambisonic soundfield manipulation uses the user’s pose (obtained from the IMU integrator) to rotate and zoom the soundfield. Then, binauralization maps the manipulated soundfield to the available speakers (we assume headphones, as is common in XR). Specifically, it applies Head-Related Transfer Functions (HRTFs) [27], which are digital filters to capture the modification of incoming sound by the person’s nose, head, and outer ear shape. Adding such cues allows the digitally rendered sound to mimic real sound, helping the brain localize sounds akin to real life. There are separate HRTFs for the left and right ears due to the asymmetrical shape of the human head.

3.4 Runtime

Figure 2 shows an ideal execution timeline for the different XR components above and their temporal dependencies. On the left, each colored rectangle boundary represents the period of the respective component — ideally the component would finish execution before the time for its next invocation. The right side shows a static representation of the dependencies among the components, illustrating the interaction between the different pipelines.

We say a consumer component exhibits a *synchronous* dependence on a producer component if the former has to wait for the last invocation of the latter to complete. An *asynchronous* dependence is softer, where the consumer can start execution with data from a previously complete invocation of the producer component. For example, the application is asynchronously dependent on the IMU integrator and VIO is synchronously dependent on the camera. The arrows in Figure 2 illustrate these inter-component dependencies, with solid indicating synchronous dependencies and dashed asynchronous.

An XR system is unlikely to follow the idealized schedule in Figure 2 (left) due to shared and constrained resources and variable running times. Thus, an explicit runtime is needed

for effective resource management and scheduling while maintaining the inter-component dependencies, resource constraints, and quality of experience.

The ILLIXR runtime schedules resources while enforcing dependencies, in part deferring to the underlying Linux scheduler. Currently, asynchronous components are scheduled by launching persistent threads that wake up at desired intervals (thread wakes are implemented via `sleep` calls). Synchronous components are scheduled whenever a new input arrives. We do not preempt invocations of synchronous components that miss their target deadline (e.g., VIO will process all camera frames that arrive). However, asynchronous components can miss work if they do not meet their target deadline (e.g., if timewarp takes too long, the display will not be updated). We do not claim that ILLIXR schedules resources optimally, only that it does so correctly (considering dependencies), and it is an extensible platform for future research. To achieve this extensibility, ILLIXR is divided into a communication framework, a set of plugins, and a plugin-loader.

The communication framework is structured around *event streams*. Event streams support writes, asynchronous reads, and synchronous reads. Synchronous reads allow a consumer to see every value produced by the producer, while asynchronous reads allow a consumer to ask for the latest value.

Plugins are shared-object files to comply with a small interface. They can only interact with other plugins through the provided communication framework. This architecture allows for modularity. Each of the components described above is incorporated into its own plugin.

The plugin loader can load a list of plugins defined at runtime. A plugin is completely interchangeable with another as long as it writes to the same event streams in the same manner. Future researchers can test alternative implementations of a single plugin without needing to reinvent the rest of the system. Development can iterate quickly, because plugins are compiled independently and linked dynamically.

ILLIXR can plug into the MonoXR OpenXR runtime [77] to support OpenXR applications. While not every OpenXR feature is implemented, we support pose queries and frame submissions, which are sufficient for the rest of the components to function. This allows ILLIXR to operate with many OpenXR-based applications, including those developed using game engines such as Unity, Unreal, and Godot (currently only Godot has OpenXR support on Linux, our supported platform).

3.5 Metrics

The ILLIXR framework provides a multitude of metrics to evaluate the quality of the system. In addition to reporting standard performance metrics such as per-component frame rate and execution time, ILLIXR reports several quality metrics: 1) motion-to-photon latency [118], a standard measure of the lag between user motion and image updates; 2) Structural Similarity Index Measure (SSIM) [119], one of the most commonly used image quality metrics in XR studies [56, 75, 101]; 3) FLIP [3], a recently introduced image quality metric that addresses the shortcomings of SSIM; and 4) pose error [130], a measure of the accuracy of the poses used to render the displayed images. The data collection in-

Table 8: Task breakdown of audio playback.

Task	Time	Sub-task	Computation	Memory Pattern
Rotation Soundfield rotation using pose	35%	Psychoacoustic filter Applies frequency-domain filter HOA rotation Rotates virtual channels	FFT; frequency domain convolution; IFFT Transcendentals; FMADDS	Butterfly pattern for FFT/IFFT; dense row-major sequential accesses for convolution Sparse column-major accesses; some temporal locality
Zoom Soundfield zoom using pose	5%	—	Linear algebra	Dense column-major sequential accesses
Binauralization HRTF application	60%	—	Identical to psychoacoustic filter	Identical to psychoacoustic filter

frastructure of ILLIXR is generic and extensible, and allows the usage of any arbitrary image quality metric. We do not yet compute a quality metric for audio beyond bitrate, but plan to add the recently developed AMBIQUAL [81] in the future.

Overall, we believe the runtime described in this paper gives researchers a jumping-point to start answering questions related to scheduling and other design decisions that trade off various QoE metrics with conventional system constraints such as PPA.

4. EXPERIMENTAL METHODOLOGY

The goals of our experiments are to quantitatively show that (1) XR presents a rich opportunity for computer architecture and systems research in domain-specific edge systems; (2) fully exploiting this opportunity requires an end-to-end system that models the complex, interacting pipelines in an XR workflow, and (3) ILLIXR is a unique testbed that provides such an end-to-end system, enabling new research directions in domain-specific edge systems.

Towards the above goals, we perform a comprehensive characterization of a live end-to-end ILLIXR system on multiple hardware platforms with representative XR workloads. We also characterize standalone components of ILLIXR using appropriate component-specific off-the-shelf datasets. This section details our experimental setup.

4.1 Experimental Setup

There are no existing systems that meet all the aspirational performance, power, and quality criteria for a fully mobile XR experience as summarized in Table 1. For our characterization, we chose to run ILLIXR on the following two hardware platforms, with three total configurations, representing a broad spectrum of power-performance tradeoffs and current XR devices (the CUDA and GLSL parts of our components as shown in Table 2 run on the GPU and the rest run on the CPU).

A high-end desktop platform: We used a state-of-the-art desktop system with an Intel Xeon E-2236 CPU (6C12T) and a discrete NVIDIA RTX 2080 GPU. This platform’s thermal design power (TDP) rating is far above what is deemed acceptable for a mobile XR system, but it is representative of the platforms on which current tethered systems run (e.g., HTC Vive Pro as in Table 1) and can be viewed as an approximate upper bound for performance of CPU+GPU based near-future embedded (mobile) systems.

An embedded platform with two configurations: We used an NVIDIA Jetson AGX Xavier development board [84] consisting of an Arm CPU (8C8T) and an NVIDIA Volta GPU. All experiments were run with the Jetson in 10 Watt

mode, the lowest possible preset. We used two different configurations – a high performance one (*Jetson-HP*) and a low power one (*Jetson-LP*). We used the maximum available clock frequencies for Jetson-HP and half those for Jetson-LP. These configurations approximate the hardware and/or the TDP rating of several commercial mobile XR devices. For example, Magic Leap One [57] used a Jetson TX2 [85], which is similar in design and TDP to our Xavier configuration. HoloLens 2 [41, 46, 105] and the Qualcomm Snapdragon 835 [1] used in Oculus Quest [91] have TDPs in the same range as our two Jetson configurations.

Our I/O setup is as follows. For the perception pipeline, we connected a ZED Mini camera [106] to the above platforms via a USB-C cable. A user walked in our lab with the camera, providing live camera and IMU input (Table 2) to ILLIXR. For the visual pipeline, we run representative VR and AR applications on a game engine on ILLIXR (Section 4.3) – these applications interact with the perception pipeline to provide the visual pipeline with the image frames to display.

We currently display the (corrected and reprojected) images on a desktop LCD monitor connected to the above hardware platforms. We use a desktop monitor due its ability to provide multiple resolution levels and refresh rates for experimentation. Although this means that the user does not see the display while walking with the camera, we practiced a trajectory that provides a reasonable response to the displayed images and used that (live) trajectory to collect results.⁴ Finally, for the audio pipeline, we use pre-recorded input (see Section 4.2).

We run our experiments for approximately 30 seconds. There is some variability due to the fact that these are real user experiments, and the user may perform slightly different actions based on application content.

4.2 Integrated ILLIXR System Configuration

The end-to-end integrated ILLIXR configuration in our experiments uses the components as described in Table 2 except for scene reconstruction, eye tracking, and hologram. The OpenXR standard currently does not support an interface for an application to use the results of scene reconstruction and only recently added an interface for eye tracking. We, therefore, do not have any applications available to use these components in an integrated setting. Although we can generate holograms, we do not yet have a holographic display setup with SLMs and other optical elements, and there are no general purpose pre-assembled off-the-shelf holographic displays available. We do report results in standalone mode for these components using off-the-shelf component-specific datasets.

⁴We plan to support North Star [79] to provide an HMD to the user.

Table 9: Key ILLIXR parameters that required manual system-level tuning. Several other parameters were tuned at the component level.

Component	Parameter Range	Tuned	Deadline
Camera (VIO)	Frame rate = 15 – 100 Hz	15 Hz	66.7 ms
	Resolution = VGA – 2K	VGA	–
	Exposure = 0.2 – 20 ms	1 ms	–
IMU (Integrator)	Frame rate = <800 Hz	500 Hz	2 ms
Display (Visual pipeline, Application)	Frame rate = 30 – 144 Hz	120 Hz	8.33 ms
	Resolution = <2K	2K	–
	Field-of-view = < 180°	90°	–
Audio (Encoding, Playback)	Frame rate = 48 – 96 Hz	48 Hz	20.8 ms
	Block size = 256 – 2048	1024	–

Configuring an XR system requires tuning multiple parameters of the different components to provide the best end-to-end user experience on the deployed hardware. This is a complex process involving the simultaneous optimization of many QoE metrics and system parameters. Currently tuning such parameters is a manual, mostly ad hoc process. Table 9 summarizes the key parameters for ILLIXR, the range available in our system for these parameters, and the final value we chose at the end of our manual tuning. We made initial guesses based on our intuition, and then chose final values based on both our perception of the smoothness of the system and profiling results. We expect the availability of ILLIXR will enable new research in more systematic techniques for performing such end-to-end system optimization.

4.3 Applications

To evaluate ILLIXR, we used four different XR applications: Sponza [78], Materials [34], Platformer [35], and a custom AR demo application with sparse graphics. The AR application contains a single light source, a few virtual objects, and an animated ball. We had to develop our own AR application because existing applications in the XR ecosystem all predominantly target VR (outside of Microsoft HoloLens). The applications were chosen for diversity of rendering complexity, with Sponza being the most graphics-intensive and AR demo being the least. We did not evaluate an MR application as OpenXR does not yet contain a scene reconstruction interface.

ILLIXR is currently only available on Linux, and since Unity and Unreal do not yet have OpenXR support on Linux, all four applications use the Godot game engine [64], a popular open-source alternative with Linux OpenXR support.

4.4 Experiments with Standalone Components

For more detailed insight into the individual components of ILLIXR, we also ran each component by itself on our desktop and embedded platforms using off-the-shelf standalone datasets. With the exception of VIO and scene reconstruction, the computation for all other components is independent of the input; therefore, the specific choice of the input dataset is not pertinent. Thus, any (monochrome) eye image for eye tracking, any frame for reprojection and hologram, and any sound sample for audio encoding and playback will produce the same result for power and performance standalone. In our experiments, we used the following datasets for these components: OpenEDS [32] for eye tracking, 2560×1440 pixel frames from VR Museum of Fine Art [102] for time-warp and hologram, and 48 KHz audio clips [29, 30] from

Freesound [28] for audio encoding and playback. For VIO and scene reconstruction, many off-the-shelf datasets are available [12, 38, 95, 108]. We used Vicon Room 1 Easy, Medium, and Difficult [12] for VIO. These sequences present varying levels of pose estimation difficulty to the VIO algorithm by varying both camera motion and level of detail in the frames. We used dyson_lab [21] for scene reconstruction as it was provided by the authors of ElasticFusion.

4.5 Metrics

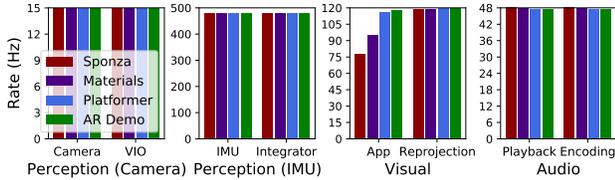
Execution time: We used NSight Systems [88] to obtain the overall execution timeline of the components and ILLIXR, including serial CPU, parallel CPU, GPU compute, and GPU graphics phases. VTune [48] (desktop only) and perf [80] provided CPU hotspot analysis and hardware performance counter information. NSight Compute [86] and Graphics [87] provided detailed information on CUDA kernels and GLSL shaders, respectively. We also developed a logging framework that allows ILLIXR to easily collect the wall time and CPU time of each of its components. We used this framework to obtain execution time without profiling overhead. We determined the contribution of a given component towards the CPU by computing the total CPU cycles consumed by that component as a fraction of the cycles used by all components.

Power and energy: For the desktop, we measured CPU power using perf and GPU power using nvidia-smi [89]. 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: CPU, GPU, DDR (DRAM), SoC (on-chip microcontrollers), and Sys (display, storage, I/O).

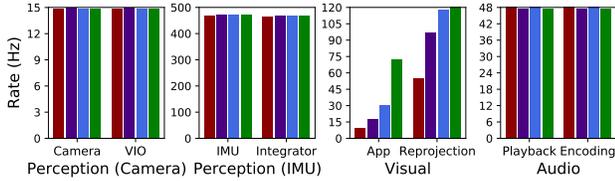
Motion-to-photon latency: We compute motion-to-photon latency as the age of the reprojected image’s pose when the pixels started to appear on screen. This latency is the sum of how old the IMU sample used for pose calculation was, the time taken by reprojection itself, and the wait time until the pixels started (but not finished) to get displayed. Mathematically, this can be formulated as: $latency = t_{imu_age} + t_{reprojection} + t_{swap}$. We do not include $t_{display}$, the time taken for all the pixels to appear on screen, in this calculation. This calculation is performed and logged by the reprojection component every time it runs. If reprojection misses vsync, the additional latency is captured in t_{swap} .

Image Quality: To compute image quality, we compare the outputs of an actual and idealized XR system. We use the Vicon Room 1 Medium dataset, which provides ground-truth poses in addition to IMU and camera data. We feed the IMU and camera data into the actual system, while feeding in the equivalent ground-truth to the idealized system. In each system, we record each image submitted by the application as well as its timestamp and pose. Due to the fact that the coordinate systems of OpenVINS and the ground truth dataset are different, we perform a trajectory alignment step to obtain aligned ground-truth poses [130]. We then feed the aligned poses back into our system to obtain the correct ground-truth images.

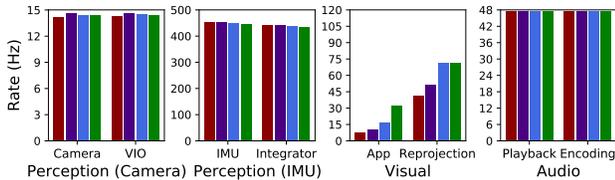
Once the images have been collected, we manually pick the starting points of the actual and idealized images sequences so that both sequences begin from the same moment in time. Then, we compute the actual and idealized reprojected images



(a) Desktop



(b) Jetson-HP



(c) Jetson-LP

Figure 3: Average frame rate for each component in the different pipelines on each application and hardware platform. The y-axis is capped at the target frame rate of the pipeline.

using the collected poses and source images. The reprojected images are the ones that would have been displayed to the user. If no image is available for a timestamp due to a dropped frame, we use the previous image. Finally, we compare the actual and idealized warped images to compute both SSIM and FLIP. We report 1-FLIP in this paper in order to be consistent with SSIM (0 being no similarity and 1 being identical images).

Pose Error: We use the data collection mechanism described above to collect actual and ground-truth poses. We calculate Absolute Trajectory Error (ATE) [108] and Relative Pose Error (RPE) [55] using OpenVINS’ error evaluation tool [92].

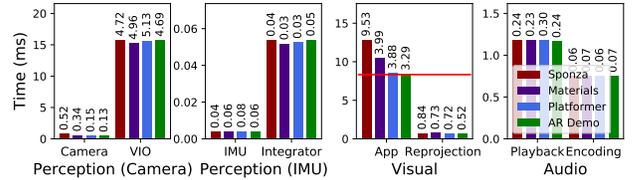
5. RESULTS

5.1 Integrated ILLIXR System

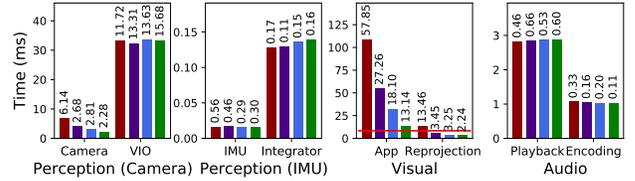
Section 5.1.1–5.1.3 respectively present system-level results for performance (execution time and throughput), power, and QoE for the integrated ILLIXR system configuration.

5.1.1 Performance

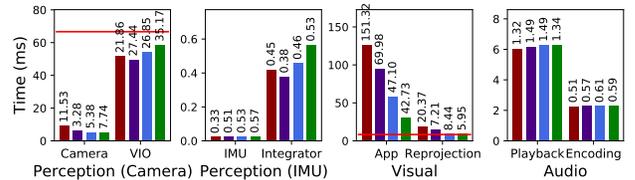
ILLIXR consists of multiple interacting streaming pipelines and there is no single number that can capture the entirety of its performance (execution time or throughput). We present the performance results for ILLIXR in terms of achieved frame rate and per-frame execution time (mean and standard deviation) for each component, and compare them against the target frame rate and corresponding deadline respectively (Table 9). Further, to understand the importance of each



(a) Desktop



(b) Jetson-HP



(c) Jetson-LP

Figure 4: Average per-frame execution time for each component for each application and hardware platform. The number on top of each bar is the standard deviation across all frames. The red horizontal line shows the maximum execution time (deadline) to meet the target frame rate. The line is not visible in graphs where the achieved time is significantly below the deadline.

component to execution time resources, we also present the contribution of the different components to the total CPU execution cycles.

Component frame rates. Figures 3(a)-(c) show each component’s average frame rate for each application, for a given hardware configuration. Components with the same target frame rate are presented in the same graph, with the maximum value on the y-axis representing this target frame rate. Thus, we show separate graphs for components in the perception, visual, and audio pipelines, with the perception pipeline further divided into two graphs representing the camera and IMU driven components.

Focusing on the desktop, Figure 3a shows that virtually all components meet, or almost meet, their target frame rates (the application component for Sponza and Materials are the only exceptions). The IMU components are slightly lower than the target frame rate only because of scheduling non-determinism at the 2 ms period required by these components. This high performance, however, comes with a significant power cost.

Moving to the lower power Jetson, we find more components missing their target frame rates. With Jetson-LP, only the audio pipeline is able to meet its target. The visual pipeline components – application and reprojection – are both severely degraded in all cases for Jetson-LP and most cases for Jetson-HP, sometimes by more than an order of magnitude. (Recall that the application and reprojection missing

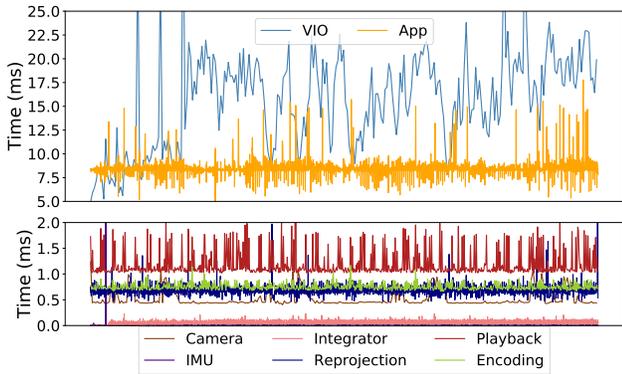


Figure 5: Per-frame execution times for Platformer on desktop. The top graph shows VIO and the application, and the bottom graph shows the remaining components. Note the different scales of the y-axes.

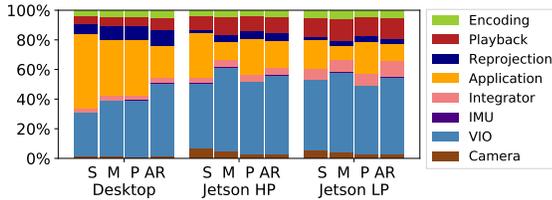


Figure 6: Contributions of ILLIXR components to CPU time.

vsync results in the loss of an entire frame. In contrast, other components can catch up on the next frame even after missing the deadline on the current frame, exhibiting higher effective frame rates.)

Although we assume modern display resolutions and refresh rates, future systems will support larger and faster displays with larger field-of-view and will integrate more components, further stressing the entire system.

Execution time per frame. Figures 4(a)-(c) show the average execution time (wall clock time) per frame for a given component, organized similar to Figure 3. The horizontal line on each graph shows the target execution time or deadline, which is the reciprocal of the target frame rate. Note that the achieved frame rate in Figure 3 cannot exceed the target frame rate as it is controlled by the runtime. The execution time per frame, however, can be arbitrarily low or high (because we don't preempt an execution that misses its deadline). The number at the top of each bar shows the standard deviation across all frames.

The mean execution times follow trends similar to those for frame rates. The standard deviations for execution time are surprisingly significant in many cases. For a more detailed view, Figure 5 shows the execution time of each frame of each ILLIXR component during the execution of Platformer on the desktop (other timelines are omitted for space). The components are split across two graphs for clarity. We expect variability in VIO (blue) and application (yellow) since these computations are known to be input-dependent. However, we see significant variability in the other components as well. This variability is due to scheduling and resource contention, motivating the need for better resource allocation, partitioning, and scheduling. A consequence of the variability is that although the mean per-frame execution time for some compo-

nents is comfortably within the target deadline, several frames do miss their deadlines; e.g., VIO in Jetson-LP. Whether this affects the user's QoE depends on how well the rest of the system compensates for these missed deadlines. For example, even if the VIO runs a little behind the camera, the IMU part of the perception pipeline may be able to compensate. Similarly, if the application misses some deadlines, reprojection may be able to compensate. Section 5.1.3 provides results for system-level QoE metrics that address these questions.

Distribution of cycles. Figure 6 shows the relative attribution of the total cycles consumed in the CPU to the different ILLIXR components for the different applications and hardware platforms. These figures do not indicate wall clock time since one elapsed cycle with two busy concurrent threads will contribute two cycles in these figures. Thus, these figures indicate total CPU resources used. We do not show GPU cycle distribution because the GPU timers significantly perturbed the rest of the execution – most of the GPU cycles are consumed by the application with the rest used by reprojection.

Focusing first on the desktop, Figure 6 shows that VIO and the application are the largest contributors to CPU cycles, with one or the other dominating, depending on the application. Reprojection and audio playback follow next, becoming larger relative contributors as the application complexity reduces. Although reprojection does not exceed 10% of the total cycles, it is a dominant contributor to motion-to-photon latency and, as discussed below, cannot be neglected for optimization.

Jetson-HP and Jetson-LP show similar trends except that we find that the application's and reprojection's relative contribution decreases while other components such as the IMU integrator become more significant relative to the desktop. This phenomenon occurs because with the more resource-constrained Jetson configurations, the application and reprojection often miss their deadline and are forced to skip the next frame (Section 3.4). Thus, the overall work performed by these components reduces, but shows up as poorer end-to-end QoE metrics discussed later (Section 5.1.3).

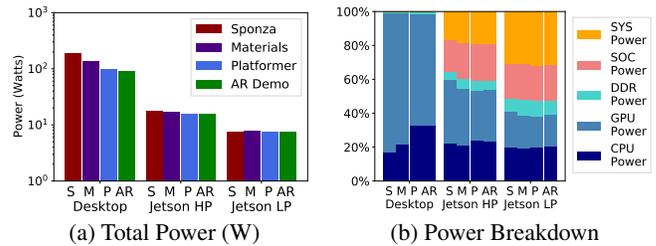


Figure 7: (a) Total power (note log scale) and (b) relative contribution to power by different hardware units for each application and hardware platform.

5.1.2 Power

Total Power. Figure 7a shows the total power consumed by ILLIXR running each application on each hardware platform. The power gap from the ideal (Table 1) is severe on all three platforms. As shown in Figure 7a, Jetson-LP, the lowest power platform, is two orders of magnitude off in terms of

Table 10: Motion-to-photon latency in milliseconds (mean \pm std dev), without $t_{display}$. Target is 20 ms for VR and 5 ms for AR (Table 1).

Application	Desktop	Jetson-hp	Jetson-lp
Sponza	3.1 \pm 1.1	13.5 \pm 10.7	19.3 \pm 14.5
Materials	3.1 \pm 1.0	7.7 \pm 2.7	16.4 \pm 4.9
Platformer	3.0 \pm 0.9	6.0 \pm 1.9	11.3 \pm 4.7
AR Demo	3.0 \pm 0.9	5.6 \pm 1.4	12.0 \pm 3.4

the ideal power while the desktop is off by three. As with performance, larger resolutions, frame rates, and field-of-views, and more components would further widen this gap.

Contribution to power from different hardware components. Figure 7b shows the relative contribution of different hardware units to the total power, broken down as CPU, GPU, DDR, SoC, and Sys. These results clearly highlight the need for studying all aspects of the system. Although the GPU dominates power on the desktop, that is not the case on Jetson. SoC and Sys power are often ignored, but doing so does not capture the true power consumption of XR workloads, which typically exercise all the aforementioned hardware units – SoC and Sys consume more than 50% of total power on Jetson-LP. Thus, reducing the power gap for XR components would require optimizing system-level hardware components as well.

5.1.3 Quality of Experience

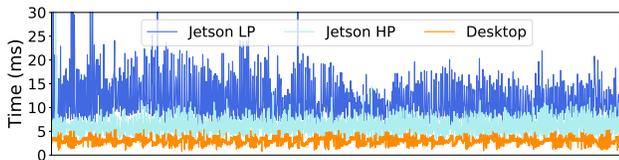


Figure 8: Motion-to-photon latency per frame of Platformer.

Quality of an XR experience is often determined through user studies [4, 6, 132]; however, these can be expensive, time consuming, and subjective. ILLIXR, therefore, provides several quantitative metrics to measure QoE. We report both results of visual examination and quantitative metrics below.

Visual examination. A detailed user study is outside the scope of this work, but there were several artifacts in the displayed image that were clearly visible. As indicated by the performance metrics, the desktop displayed smooth images for all four applications. Jetson-HP showed perceptibly increased judder for Sponza. Jetson-LP showed dramatic pose drift and clearly unacceptable images for Sponza and Materials, though it was somewhat acceptable for the less intense Platformer and AR Demo. As discussed in Section 5.1.1, the average VIO frame rate for Jetson-LP stayed high, but the variability in the per-frame execution time resulted in many missed deadlines, which could not be fully compensated by the IMU or reprojection. These effects are quantified with the metrics below.

Motion-to-photon latency (MTP). Table 10 shows the mean and standard deviation of MTP for all cases. Figure 8 shows MTP for each frame over the execution of Platformer on all hardware (we omit the other applications for space). Recall the target MTP is 20 ms for VR and 5 ms for AR (Table 1) and that our measured numbers reported do not

Table 11: Image and pose quality metrics (mean \pm std dev) for Sponza.

Platform	SSIM	1-FLIP	ATE/degree	ATE/meters
Desktop	0.83 \pm 0.04	0.86 \pm 0.05	8.6 \pm 6.2	0.33 \pm 0.15
Jetson-hp	0.80 \pm 0.05	0.85 \pm 0.05	18 \pm 13	0.70 \pm 0.33
Jetson-lp	0.68 \pm 0.09	0.65 \pm 0.17	138 \pm 26	13 \pm 10

include $t_{display}$ (4.2 ms).

Table 10 and our detailed per-frame data shows that the desktop can achieve the target VR MTP for virtually all frames. For AR, most cases appear to meet the target, but adding $t_{display}$ significantly exceeds the target. Both Jetsons cannot make the target AR MTP for an average frame. Jetson-HP is able to make the VR target MTP for the average frame for all applications and for most frames for all except Sponza (even after adding $t_{display}$). Jetson-LP shows a significant MTP degradation – on average, it still meets the target VR MTP, but both the mean and variability increase with increasing complexity of the application until Sponza is practically unusable (19.3 ms average without $t_{display}$ and 14.5 ms standard deviation). Thus, the MTP data collected by ILLIXR is consistent with the visual observations reported above.

Offline metrics for image and pose quality. Table 11 shows the mean and standard deviation for SSIM, 1-FLIP, ATE/degree, and ATE/distance (Section 3.5 and 4.5) for Sponza on all hardware configurations. (We omit other applications for space.) Recall that these metrics require offline analysis and use a different offline trajectory dataset for VIO (that comes with ground truth) from the live camera trajectory reported in the rest of this section. Nevertheless, the visual experience of the applications is similar. We find that all metrics degrade as the hardware platform becomes more constrained. While the trajectory errors reported are clearly consistent with the visual experience, the SSIM and FLIP values seem deceptively high for the Jetsons (specifically the Jetson-LP where VIO shows a dramatic drift). Quantifying image quality for XR is known to be challenging [50, 129]. While some work has proposed the use of more conventional graphics-inspired metrics such as SSIM and FLIP, this is still a topic of ongoing research [63, 128]. ILLIXR is able to capture the proposed metrics, but our work also motivates and enables research on better image quality metrics for XR experiences.

5.1.4 Key Implications for Architects

Our results characterize end-to-end performance, power, and QoE of an XR device, exposing new research opportunities for architects and demonstrating ILLIXR as a unique testbed to enable exploration of these opportunities, as follows.

Performance, power, and QoE gaps. Our results quantitatively show that collectively there is several orders of magnitude performance, power, and QoE gap between current representative desktop and embedded class systems and the goals in Table 1. The above gap will be further exacerbated with higher fidelity displays and the addition of more components for a more feature-rich XR experience (e.g., scene reconstruction, eye tracking, hand tracking, and holography).

While the presence of these gaps itself is not a surprise, we provide the first such quantification and analysis. This provides insights for directions for architecture and systems

research (below) as well as demonstrates ILLIXR as a one-of-a-kind testbed that can enable such research.

No dominant performance component. Our more detailed results show that there is no one component that dominates all the metrics of interest. While the application and VIO seem to dominate CPU cycles, reprojection latency is critical to MTP, and audio playback takes similar or more cycles as reprojection. Since image quality relies on accurate poses, frequent operation of the IMU integrator is essential to potentially compensate for missed VIO deadlines. As mentioned earlier, Hologram dominated the GPU and was not even included in the integrated configuration studied since it precluded meaningful results on the Jetson-LP. Thus, to close the aforementioned gaps, *all* components have to be considered together, even those that may appear relatively inexpensive at first glance. Moreover, we expect this diversity of important components to only increase as more components are included for more feature rich experiences. These observations coupled with the large performance-power gaps above indicate a rich space for hardware specialization research, including research in automated tool flows to determine what to accelerate, how to accelerate these complex and diverse codes, and how to best exploit accelerator-level parallelism [40].

Full-system power contributions. Our results show that addressing the power gap requires considering *system-level* hardware components, such as display and other I/O, which consume significant energy themselves. This also includes numerous sensors. While we do not measure individual sensor power, it is included in Sys power on Jetson, which is significant. This motivates research in unconventional architecture paradigms such as on-sensor computing to save I/O power.

Variability and scheduling. Our results show large variability in per-frame processing times in many cases, either due to inherent input-dependent nature of the component (e.g., VIO) or due to resource contention from other components. This variability poses challenges to, and motivates research directions in, scheduling and resource partitioning and allocation of hardware resources. As presented in Table 9, there are a multitude of parameters that need to be tuned for optimal performance of the system. This paper performs manual tuning to fix these parameters; however, ideally they would adapt to the performance variations over time, co-optimized with scheduling and resource allocation decisions designed to meet real-time deadlines within power constraints to provide maximal QoE. ILLIXR provides a flexible testbed to enable such research.

End-to-end quality metrics. The end-to-end nature of ILLIXR allows it to provide end-to-end quality metrics such as MTP and see the impact of optimizations on the final displayed images. Our results show that per-component metrics (e.g., VIO frame rate) are insufficient to determine the impact on the final user experience (e.g., as in Jetson-LP for Sponza). At the same time, there is a continuum of acceptable experiences (as in Sponza for desktop and Jetson-HP). Techniques such as approximate computing take advantage of such applications but often focus on subsystems that make it hard to assess end user impact. ILLIXR provides a unique testbed to develop cross-layer approximate computing techniques that

can be driven by end-user experience. Our results also motivate work on defining more perceptive quantitative metrics for image quality and ILLIXR provides a testbed for such research.

5.2 Per Component Analysis

The results so far motivate specialization of ILLIXR components to meet the performance-power-QoE gap as seen through the end-to-end characterization. This section examines the components in isolation to provide insights on how to specialize. We performed an extensive deep dive analysis of all the components. Tables 3–8 summarize the tasks within each component, including the key computation and memory patterns. Figure 10 illustrates the contribution of the tasks to a component’s execution time. Figure 9 shows the CPU IPC and cycles breakdown for each component in aggregate. Space restrictions prevent a detailed deep dive for each component, but we present two brief case studies.

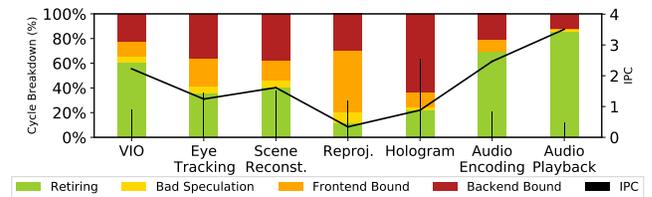


Figure 9: Cycle breakdown and IPC of ILLIXR components.

Component Diversity The components present unique characteristics. Figure 9 shows a wide range in both IPC and hardware utilization for the different components. Figure 10 shows that the tasks are typically unique to a component (Section 3).

Task Diversity There is a remarkable diversity of algorithmic tasks, ranging from feature detection to map fusion to signal processing and more. These tasks are differently amenable for execution on the CPU, GPU-compute, or GPU-graphics.

Task Dominance No component is composed of just one task. The most homogeneous is audio encoding where ambisonic encoding is 81% of the total – but accelerating just that would limit the overall speedup to $5.3\times$ by Amdahl’s law, which may not be sufficient given the power and performance gap shown in Section 5.1. VIO is the most diverse, with seven major tasks and many more sub-tasks.

Input-Dependence Although we saw significant per-frame execution time variability in all components in the integrated system, the only components that exhibit such variability standalone are VIO and scene reconstruction. VIO shows a range of execution time throughout its execution, with a coefficient of variation from 17% to 26% across the studied datasets [12]. Scene reconstruction’s execution time keeps steadily increasing due to the increasing size of its map. Loop closure attempts result in execution time spikes of 100’s of ms, an order of magnitude more than its average per-frame execution time.

Case Study: VIO OpenVINS consists of 7 tasks, each with their own sub-tasks and unique characteristics. The compute patterns span stencils, linear algebra, Gauss-Newton, and QR decomposition, among others. The memory patterns are

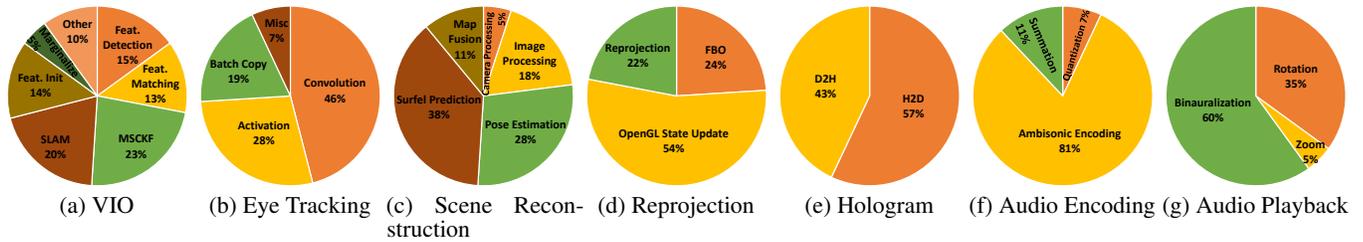


Figure 10: Execution time breakdown into tasks of key components on the desktop. There is a large diversity in tasks and no task dominates.

equally diverse, spanning dense and sparse, local and global, and row-major and column-major accesses. These properties manifest themselves in microarchitectural utilization as well. For instance, the core task of feature matching, KLT, has an IPC of 3, which is significantly higher than the component average of 2.2 due to its computationally intensive integer stencil. These properties suggest that most sub-tasks may require separate accelerators. Furthermore, these accelerators must be designed to accommodate for significant variations in execution time.

Case Study: Audio Playback Audio playback is an example of a component for which building a unified accelerator is straightforward. Psychoacoustic filter and binauralization have **identical** compute and memory characteristics, while HOA rotation and zoom have similar (but not identical) characteristics. Consequently, a traditional audio accelerator with systolic arrays, specialized FFT blocks, and streambuffers should suffice for all four sub-tasks.

Key implications for architects The diversity of components and tasks, the lack of a single dominant task for most components, and per-frame variability pose several challenges for hardware specialization. It is likely impractical to build a unique accelerator for every task given the large number of tasks and the severe power and area constraints for XR devices: leakage power will be additive across accelerators, and interfaces between these accelerators and other peripheral logic will further add to this power and area (we identified 27 tasks in Figure 10, and expect more tasks with new components).

This motivates *automated techniques for determining what to accelerate*. An open question is what is the granularity of acceleration; e.g., one accelerator per task or a more general accelerator that supports *common primitives shared among components*. Besides the design of the accelerators themselves, architects must also determine *accelerator communication interfaces* and techniques to exploit accelerator level parallelism.

Finally, although not explicitly studied here, algorithms for the various components continue to evolve and there are multiple choices available at any time (ILLIXR already supports such choices). This motivates *programmable accelerators* with the attendant challenges of developing a software stack for such a system. ILLIXR provides a testbed to perform all of the above research in the context of the full system.

6. 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: Processing-in-Memory architectures for graphics [125, 126], video accelerators for VR [58, 59, 70, 71], 3D point cloud acceleration [24, 127], stereo acceleration [25], computer vision accelerators [112, 133], and SLAM chips [69, 109]. However, these works have been for single components in isolation. Significant work has also been done on remote rendering [9, 61, 66, 67, 75, 131] and 360 video streaming [5, 20, 43, 96]. While these works do consider full end-to-end systems, they only focus on the rendering aspect of the system and do not consider the interplay between the various components of the system. Our work instead provides a full-system-benchmark, consisting of a set of XR components representing a full XR workflow, along with insights into their characteristics. Our goal is to propel architecture research in the direction of full domain-specific *system* design.

There have been other works that have developed benchmark suites for XR. Raytrace in SPLASH-2 [124], VRMark [117], FCAT [90], and Unigine [114] are examples of graphics rendering benchmarks, but do not contain the perception pipeline nor adaptive display components. The SLAMBench [8, 11, 82] 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.

Chen et al. [17] characterized AR applications on a commodity smartphone, but neither analyzed individual components nor considered futuristic components.

7. CONCLUSION

We developed ILLIXR, the first open source end-to-end XR testbed to democratize XR systems research. Our analysis enabled by ILLIXR shows a number of interesting results for architects – demanding performance and power requirements; a large diversity of critical tasks; significant input dependent computation that challenges scheduling and specialization; and a diversity in bottlenecks throughout the system, from graphics and compute requirements to power and resource contention. ILLIXR and our analysis have the potential to propel many new directions in architecture research. Some research avenues enabled include: generation of high-level, programmable, specialized compute and memory primitives for shared specialized hardware; SoC design with accelerator level parallelism; inter-accelerator communication and coherence architectures; programming languages, IRs, and

OSes for domain-specific systems; real-time scheduling and resource management of domain-specific SoCs; QoE metrics research (including approximate computing); on-sensor computing; device-edge-cloud partitioning of workloads; and full-system simulator development. We, in collaboration with others, are already pursuing several of these directions. Although ILLIXR already incorporates a representative 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 will continue to evolve, providing an increasingly comprehensive resource to solve XR systems' most difficult challenges.

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 Cui, Aleksandra Faust, Liang Gao, Matt Horsnell, Amit Jindal, Steve LaValle, Steve Lovegrove, Andrew Maimone, Vegard Øye, Martin Persson, Archontis Politis, Eric Shaffer, and Paris Smaragdiss.

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, the National Science Foundation under grant CCF 16-19245, 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] F. Abazovic, "Snapdragon 835 is a 10nm slap in Intel's face," 2018. [Online]. Available: <https://www.fudzilla.com/news/mobile/42154-snapdragon-835-in-10nm-is-slap-in-intel-s-face>
- [2] R. Aggarwal, T. P. Grantcharov, J. R. Eriksen, D. Bliurup, 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. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/16858196>
- [3] P. Andersson, T. Akenine-Möller, J. Nilsson, K. Åström, M. Oskarsson, and M. Fairchild, "Flip: A difference evaluator for alternating images," *Proceedings of the ACM in Computer Graphics and Interactive Techniques*, vol. 3, no. 2, 2020.
- [4] Y. Arifin, T. G. Sastria, and E. Barlian, "User experience metric for augmented reality application: a review," *Procedia Computer Science*, vol. 135, pp. 648–656, 2018.
- [5] Y. Bao, T. Zhang, A. Pande, H. Wu, and X. Liu, "Motion-prediction-based multicast for 360-degree video transmissions," in *2017 14th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, 2017, pp. 1–9.
- [6] B. Bauman and P. Seeling, "Towards predictions of the image quality of experience for augmented reality scenarios," *arXiv preprint arXiv:1705.01123*, 2017.
- [7] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC benchmark suite: Characterization and architectural implications," in *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques*, ser. PACT '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 72–81. [Online]. Available: <https://doi.org/10.1145/1454115.1454128>
- [8] B. Bodin, H. Wagstaff, S. Saecdi, L. Nardi, E. Vespa, J. Mawer, A. Nisbet, M. Lujan, S. Furber, A. J. Davison, P. H. J. Kelly, and M. F. P. O'Boyle, "Slambench2: Multi-objective head-to-head benchmarking for visual slam," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 3637–3644.
- [9] K. Boos, D. Chu, and E. Cuervo, "Flashback: Immersive virtual reality on mobile devices via rendering memoization," in *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 291–304. [Online]. Available: <https://doi.org/10.1145/2906388.2906418>
- [10] J. Bucek, K.-D. Lange, and J. v. Kistowski, "Spec cpu2017: Next-generation compute benchmark," in *Companion of the 2018 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 41–42. [Online]. Available: <https://doi.org/10.1145/3185768.3185771>
- [11] M. Bujanca, P. Gafton, S. Saecdi, A. Nisbet, B. Bodin, M. F. P. O'Boyle, A. J. Davison, P. H. J. Kelly, G. Riley, B. Lennox, M. Luján, and S. Furber, "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)*, 2019, pp. 6351–6358.
- [12] 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.
- [13] L. Carlone, Z. Kira, C. Beall, V. Indelman, and F. Dellaert, "Eliminating conditionally independent sets in factor graphs: A unifying perspective based on smart factors," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 4290–4297.
- [14] 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>
- [15] A. K. Chaudhary, R. Kothari, M. Acharya, S. Dangi, N. Nair, R. Bailey, C. Kanan, G. Diaz, and J. B. Pelz, "Ritnet: Real-time semantic segmentation of the eye for gaze tracking," in *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 2019, pp. 3698–3702.
- [16] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S. Lee, and K. Skadron, "Rodinia: A benchmark suite for heterogeneous computing," in *2009 IEEE International Symposium on Workload Characterization (IISWC)*, 2009, pp. 44–54.
- [17] H. Chen, Y. Dai, H. Meng, Y. Chen, and T. Li, "Understanding the Characteristics of Mobile Augmented Reality Applications," in *2018 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, April 2019, pp. 128–138.
- [18] M. L. Chénéchal and J. C. Goldman, "HTC Vive Pro Time Performance Benchmark for Scientific Research," in *ICAT-EGVE 2018 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*, G. Bruder, S. Yoshimoto, and S. Cobb, Eds. The Eurographics Association, 2018.
- [19] I. Cutress, "Spotted: Qualcomm Snapdragon 8cx wafer on 7nm," <https://www.anandtech.com/show/13687/qualcomm-snapdragon-8cx-wafer-on-7nm>, 2018.
- [20] M. Dasari, A. Bhattacharya, S. Vargas, P. Sahu, A. Balasubramanian, and S. R. Das, "Streaming 360-degree videos using super-resolution," in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020, pp. 1977–1986.
- [21] "ElasticFusion repository," <http://github.com/mp3guy/ElasticFusion/>, 2015.
- [22] M. S. Elbamby, C. Perfecto, M. Bennis, and K. Doppler, "Toward low-latency and ultra-reliable virtual reality," *IEEE Network*, vol. 32, no. 2, pp. 78–84, 2018.
- [23] 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.
- [24] Y. Feng, B. Tian, T. Xu, P. Whatmough, and Y. Zhu, "Mesorasi:

- Architecture support for point cloud analytics via delayed-aggregation,” *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Oct 2020. [Online]. Available: <http://dx.doi.org/10.1109/MICRO50266.2020.00087>
- [25] 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>
- [26] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, “Imu preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation,” *Robotics: Science and Systems (RSS) conference*, 01 2015.
- [27] 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>
- [28] “Freesound,” <https://freesound.org/>.
- [29] “Science Teacher Lecturing,” May 2013. [Online]. Available: <https://freesound.org/people/SpliceSound/sounds/188214/>
- [30] “Radio Recording,” <https://freesound.org/people/waveplay./sounds/397000/>, July 2017.
- [31] 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.
- [32] S. J. Garbin, Y. Shen, I. Schuetz, R. Cavin, G. Hughes, and S. S. Talathi, “OpenEDS: Open Eye Dataset,” 2019.
- [33] P. Geneva, K. Eckenhoff, W. Lee, Y. Yang, and G. Huang, “OpenVins: A research platform for visual-inertial estimation,” *IROS 2019 Workshop on Visual-Inertial Navigation: Challenges and Applications*, 2019.
- [34] Godot, “Material testers,” https://github.com/godotengine/godot-demo-projects/tree/master/3d/material_testers, 2020.
- [35] Godot, “Platformer 3D,” <https://github.com/godotengine/godot-demo-projects/tree/master/3d/platformer>, 2020.
- [36] L. Goode, “The hololens 2 puts a full-fledged computer on your face,” <https://www.wired.com/story/microsoft-hololens-2-headset/>, Feb 2019.
- [37] “GTSAM repository,” 2020. [Online]. Available: <https://github.com/borglab/gtsam>
- [38] A. Handa, T. Whelan, J. McDonald, and A. Davison, “A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM,” in *IEEE Intl. Conf. on Robotics and Automation, ICRA*, Hong Kong, China, May 2014.
- [39] J. L. Henning, “Spec cpu2006 benchmark descriptions,” *SIGARCH Comput. Archit. News*, vol. 34, no. 4, p. 1–17, Sep. 2006. [Online]. Available: <https://doi.org/10.1145/1186736.1186737>
- [40] M. D. Hill and V. J. Reddi, “Accelerator-level parallelism,” *CoRR*, vol. abs/1907.02064, 2019. [Online]. Available: <http://arxiv.org/abs/1907.02064>
- [41] K. Hinum, “Qualcomm Snapdragon 855 SoC - Benchmarks and Specs,” 2019. [Online]. Available: <https://www.notebookcheck.net/Qualcomm-Snapdragon-855-SoC-Benchmarks-and-Specs.375436.0.html>
- [42] F. Hollerweger, “An introduction to higher order ambisonic,” 2008.
- [43] X. Hou, S. Dey, J. Zhang, and M. Budagavi, “Predictive view generation to enable mobile 360-degree and vr experiences,” in *Proceedings of the 2018 Morning Workshop on Virtual Reality and Augmented Reality Network*, ser. VR/AR Network ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 20–26. [Online]. Available: <https://doi.org/10.1145/3229625.3229629>
- [44] HTC, “HTC VIVE Pro,” 2016. [Online]. Available: <https://www.vive.com/us/product/vive-pro/>
- [45] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” 2016.
- [46] Ian Cutress, “Hot chips 31 live blogs: Microsoft hololens 2.0 silicon,” <https://www.anandtech.com/show/14775/hot-chips-31-live-blogs-microsoft-hololens-20-silicon>, 2019.
- [47] Intel, “Intelrealsense,” <https://github.com/IntelRealSense/librealsense>, 2015.
- [48] Intel, “Intel vtune profiler,” 2020. [Online]. Available: <https://software.intel.com/content/www/us/en/develop/tools/vtune-profiler.html>
- [49] D. Kanter, “Graphics processing requirements for enabling immersive vr,” *Whitepaper*, 2015.
- [50] H. G. Kim, H.-T. Lim, and Y. M. Ro, “Deep virtual reality image quality assessment with human perception guider for omnidirectional image,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 4, pp. 917–928, 2019.
- [51] “KimeraVIO repository,” <https://github.com/MIT-SPARK/Kimera-VIO>, 2017.
- [52] “KinectFusionApp repository,” <https://github.com/chrdiller/KinectFusionApp>, 2018.
- [53] 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>
- [54] S. Krohn, J. Tromp, E. M. Quinque, J. Belger, F. Klotzsche, S. Rekers, P. Chojecki, J. de Mooij, M. Akbal, C. McCall, A. Villringer, M. Gaebler, C. Finke, and A. Thöne-Otto, “Multidimensional evaluation of virtual reality paradigms in clinical neuropsychology: Application of the vr-check framework,” *J Med Internet Res*, vol. 22, no. 4, p. e16724, Apr 2020. [Online]. Available: <https://doi.org/10.2196/16724>
- [55] R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, and A. Kleiner, “On measuring the accuracy of slam algorithms,” *Autonomous Robots*, vol. 27, no. 4, p. 387, 2009.
- [56] Z. Lai, Y. C. Hu, Y. Cui, L. Sun, N. Dai, and H.-S. Lee, “Furion: Engineering high-quality immersive virtual reality on today’s mobile devices,” *IEEE Transactions on Mobile Computing*, 2019.
- [57] M. Leap, “Magic Leap 1,” 2019. [Online]. Available: <https://www.magicleap.com/en-us/magic-leap-1>
- [58] 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>
- [59] Y. Leng, C.-C. Chen, Q. Sun, J. Huang, and Y. Zhu, “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>
- [60] 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>
- [61] Y. Li and W. Gao, “Muvr: Supporting multi-user mobile virtual reality with resource constrained edge cloud,” in *2018 IEEE/ACM Symposium on Edge Computing (SEC)*, 2018, pp. 1–16.
- [62] “libspatialaudio repository,” <https://github.com/videolabs/libspatialaudio>, 2019.
- [63] H.-T. Lim, H. G. Kim, and Y. M. Ra, “Vr iqa net: Deep virtual reality image quality assessment using adversarial learning,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6737–6741.
- [64] J. Linietsky, A. Manzur, and contributors, “Godot,” <https://godotengine.org/>, 2020.
- [65] 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. 11 562–11 573, May 2010. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-18-11-11562>
- [66] T. Liu, S. He, S. Huang, D. Tsang, L. Tang, J. Mars, and W. Wang, “A benchmarking framework for interactive 3d applications in the cloud,” *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Oct 2020. [Online]. Available: <http://dx.doi.org/10.1109/MICRO50266.2020.00076>

- [67] X. Liu, C. Vlachou, F. Qian, C. Wang, and K.-H. Kim, "Firefly: Untethered multi-user VR for commodity mobile devices," in *2020 USENIX Annual Technical Conference (USENIX ATC 20)*. USENIX Association, Jul. 2020, pp. 943–957. [Online]. Available: <https://www.usenix.org/conference/atc20/presentation/liu-xing>
- [68] 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.
- [69] 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.
- [70] 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.
- [71] 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>
- [72] 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/>
- [73] M. McGuire, "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/>
- [74] 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>
- [75] J. Meng, S. Paul, and Y. C. Hu, "Coterie: Exploiting frame similarity to enable high-quality multiplayer vr on commodity mobile devices," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 923–937. [Online]. Available: <https://doi.org/10.1145/3373376.3378516>
- [76] Microsoft, "Microsoft hololens 2," 2019. [Online]. Available: <https://www.microsoft.com/en-us/hololens/hardware>
- [77] "Monado - open source XR platform." Available at <https://monado.dev/> (accessed April 5, 2020).
- [78] Monado, "Sponza scene in Godot with OpenXR addon," <https://gitlab.freedesktop.org/monado/demos/godot-sponza-openxr>, 2019.
- [79] L. Motion, "Project North Star," 2020. [Online]. Available: <https://github.com/leapmotion/ProjectNorthStar>
- [80] N/A, "perf," 2020. [Online]. Available: https://perf.wiki.kernel.org/index.php/Main_Page
- [81] M. Narbut, J. Skoglund, A. Allen, M. Chinen, D. Barry, and A. Hines, "Ambiquat: Towards a quality metric for headphone rendered compressed ambisonic spatial audio," *Applied Sciences*, vol. 10, no. 9, 2020. [Online]. Available: <https://www.mdpi.com/2076-3417/10/9/3188>
- [82] L. Nardi, B. Bodin, M. Z. Zia, J. Mawer, A. Nisbet, P. H. J. Kelly, A. J. Davison, M. Luján, M. F. P. O'Boyle, G. Riley, N. Topham, and S. Furber, "Introducing slambench, a performance and accuracy benchmarking methodology for slam," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 5783–5790.
- [83] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *2011 10th IEEE International Symposium on Mixed and Augmented Reality*, 2011, pp. 127–136.
- [84] NVIDIA, "NVIDIA AGX Xavier," 2017. [Online]. Available: <https://developer.nvidia.com/embedded/jetson-agx-xavier-developer-kit>
- [85] NVIDIA, "NVIDIA TX2," 2017. [Online]. Available: <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-tx2/>
- [86] NVIDIA, "Nsight compute," 2020. [Online]. Available: <https://developer.nvidia.com/nsight-compute>
- [87] NVIDIA, "Nsight graphics," 2020. [Online]. Available: <https://developer.nvidia.com/nsight-graphics>
- [88] NVIDIA, "Nsight systems," 2020. [Online]. Available: <https://developer.nvidia.com/nsight-systems>
- [89] NVIDIA, "Nvidia system management interface," 2020. [Online]. Available: <https://developer.nvidia.com/nvidia-system-management-interface>
- [90] "NVIDIA Frame Capture Analysis Tool," 2017. [Online]. Available: <https://www.geforce.com/hardware/technology/fcat>
- [91] Oculus, "Oculus Quest: All-in-One VR Headset," 2019. [Online]. Available: https://www.oculus.com/quest/?locale=en_US
- [92] "OpenVINS repository," https://github.com/rpng/open_vins, 2019.
- [93] 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>
- [94] M. Persson, D. Engström, and M. Goksör, "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>
- [95] T. Pire, M. Mujica, J. Civera, and E. Kofman, "The rosario dataset: Multisensor data for localization and mapping in agricultural environments," *The International Journal of Robotics Research*, vol. 38, no. 6, pp. 633–641, 2019.
- [96] F. Qian, B. Han, Q. Xiao, and V. Gopalakrishnan, "Flare: Practical viewport-adaptive 360-degree video streaming for mobile devices," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 99–114. [Online]. Available: <https://doi.org/10.1145/3241539.3241565>
- [97] Rebecca Pool, "Ar/vr/mr 2020: The future now arriving," 2017. [Online]. Available: <http://microvision.blogspot.com/2020/02/arvr-mr-2020-future-now-arriving.html>
- [98] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.
- [99] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020. [Online]. Available: <https://github.com/MIT-SPARK/Kimera>
- [100] C. Sakalis, C. Leonardsson, S. Kaxiras, and A. Ros, "Splash-3: A properly synchronized benchmark suite for contemporary research," in *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2016, pp. 101–111.
- [101] A. Schollmeyer, S. Schneegans, S. Beck, A. Steed, and B. Froehlich, "Efficient hybrid image warping for high frame-rate stereoscopic rendering," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 4, pp. 1332–1341, 2017.
- [102] F. Sinclair, "The vr museum of fine art," 2016. [Online]. Available: https://store.steampowered.com/app/515020/The_VR_Museum_of_Fine_Art/
- [103] J. P. Singh, W.-D. Weber, and A. Gupta, "Splash: Stanford parallel applications for shared-memory," *SIGARCH Comput. Archit. News*, vol. 20, no. 1, p. 5–44, Mar. 1992. [Online]. Available: <https://doi.org/10.1145/130823.130824>
- [104] Skarredghost, "All you need to know on hololens 2," <https://skarredghost.com/2019/02/24/all-you-need-know-hololens-2/>, May 2019.

- [105] R. Smith and A. Frumusanu, "The Snapdragon 845 Performance Preview: Setting the Stage for Flagship Android 2018," <https://www.anandtech.com/show/12420/snapdragon-845-performance-preview/4>, 2018.
- [106] Stereolabs, "ZED Mini - Mixed-Reality Camera," 2018. [Online]. Available: <https://www.stereolabs.com/zed-mini/>
- [107] Stereolabs, "ZED Software Development Kit," 2020. [Online]. Available: <https://www.stereolabs.com/developers/release/>
- [108] J. Sturm, S. Magnenat, N. Engelhard, F. Pomerleau, F. Colas, W. Burgard, D. Cremers, and R. Siegwart, "Towards a benchmark for RGB-D SLAM evaluation," in *Proc. of the RGB-D Workshop on Advanced Reasoning with Depth Cameras at Robotics: Science and Systems Conf. (RSS)*, Los Angeles, USA, June 2011.
- [109] 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.
- [110] D. Takahashi, "Oculus chief scientist mike abrash still sees the rosy future through ar/vr glasses," <https://venturebeat.com/2018/09/26/oculus-chief-scientist-mike-abrash-still-sees-the-rosy-future-through-ar-vr-glasses/>, September 2018.
- [111] The Khronos Group Inc., "The openxr specification," Available at <https://www.khronos.org/registry/OpenXR/specs/1.0/html/xrspec.html> (accessed April 5, 2020), Mar. 2020, version 1.0.8.
- [112] 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.
- [113] P. R. K. Turnbull and J. R. Phillips, "Ocular effects of virtual reality headset wear in young adults," in *Scientific Reports*, 2017.
- [114] "Unigine Superposition Benchmark," 2017. [Online]. Available: <https://benchmark.unigine.com/superposition>
- [115] D. Van Krevelen and R. Poelman, "A survey of augmented reality technologies, applications and limitations," *International journal of virtual reality*, vol. 9, no. 2, pp. 1–20, 2010.
- [116] 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.
- [117] "VRMark," 2016. [Online]. Available: <https://benchmarks.ul.com/vrmark>
- [118] 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>
- [119] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [120] "van Waveren Oculus Demo repository," 2017. [Online]. Available: <https://github.com/KhronosGroup/Vulkan-Samples-Deprecated/tree/master/samples/apps/atw>
- [121] T. Whelan, S. Leutenegger, R. Moreno, B. Glocker, and A. Davison, "ElasticFusion: Dense SLAM without a pose graph," in *Robotics: Science and Systems*, 2015.
- [122] Wikipedia contributors, "Hololens 2 — Wikipedia, the free encyclopedia," https://en.wikipedia.org/wiki/HoloLens_2, 2019.
- [123] Wikipedia contributors, "Htc vive — Wikipedia, the free encyclopedia," https://en.wikipedia.org/w/index.php?title=HTC_Vive&oldid=927924955, 2019.
- [124] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The splash-2 programs: Characterization and methodological considerations," in *Proceedings of the 22nd Annual International Symposium on Computer Architecture*, ser. ISCA '95. New York, NY, USA: Association for Computing Machinery, 1995, p. 24–36. [Online]. Available: <https://doi.org/10.1145/223982.223990>
- [125] 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.
- [126] 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.
- [127] 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>
- [128] J. Yang, T. Liu, B. Jiang, H. Song, and W. Lu, "3d panoramic virtual reality video quality assessment based on 3d convolutional neural networks," *IEEE Access*, vol. 6, pp. 38 669–38 682, 2018.
- [129] B. Zhang, J. Zhao, S. Yang, Y. Zhang, J. Wang, and Z. Fei, "Subjective and objective quality assessment of panoramic videos in virtual reality environments," in *2017 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*. IEEE, 2017, pp. 163–168.
- [130] Z. Zhang and D. Scaramuzza, "A tutorial on quantitative trajectory evaluation for visual (-inertial) odometry. in 2018 ieee," in *RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7244–7251.
- [131] S. Zhao, H. Zhang, S. Bhuyan, C. S. Mishra, Z. Ying, M. T. Kandemir, A. Sivasubramaniam, and C. R. Das, "Deja View: Spatio-Temporal Compute Reuse for Energy-Efficient 360 VR Video Streaming," in *Proceedings of the 47th International Symposium on Computer Architecture*, ser. ISCA '20, 2020.
- [132] Y. Zhu, X. Min, D. Zhu, K. Gu, J. Zhou, G. Zhai, X. Yang, and W. Zhang, "Toward better understanding of saliency prediction in augmented 360 degree videos," 2020.
- [133] 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>