32ND USENIX SECURITY SYMPOSIUM

It's all in your head(set): Side-channel attacks on AR/VR systems

Yicheng Zhang, Carter Slocum, Jiasi Chen, Nael Abu-Ghazaleh

yzhan846@ucr.edu

University of California, Riverside

UC RIVERSIDE

AR/VR Systems – An Immersive Lifestyle

- Augmented Reality/Virtual Reality (AR/VR) systems are everywhere.

Microsoft HoloLens

oculus

Vision Pro

VIVE

AR/VR Systems – An Immersive Lifestyle

- Augmented Reality/Virtual Reality (AR/VR) systems are everywhere.
- Over 60 Billion US Dollar Market Size.

Microso
HoloLen

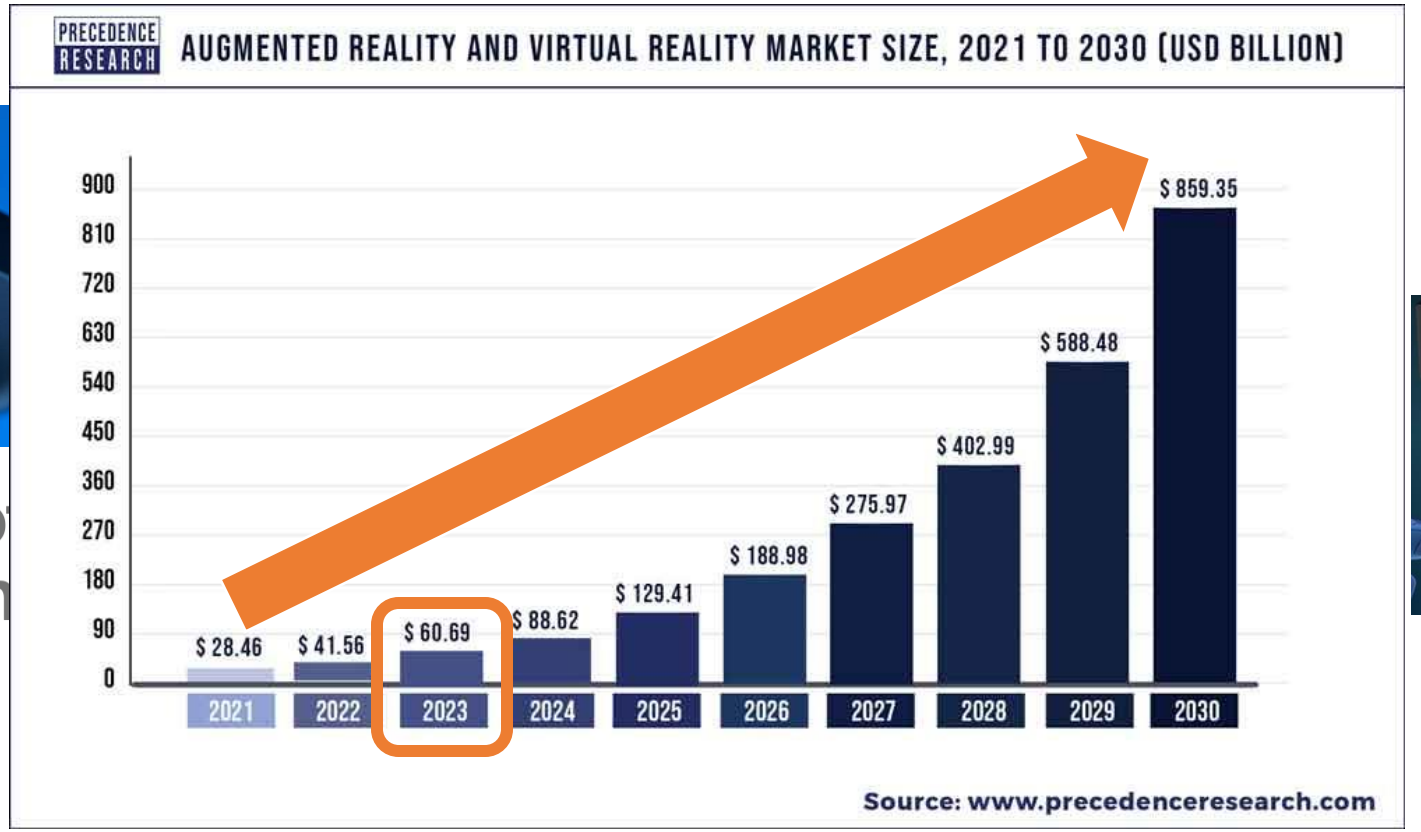

VIVE

Outline

- Background: Software and Hardware Architecture for AR/VR Systems.
- Threat model and Leakage Vectors.

- Three Classes of Side-channel Attacks.
  - User Interaction: Hand gestures, Voice input, and Keystrokes.
  - Concurrent Application: Application Launch.
  - Real-world Environment: Bystander Estimation.

- Mitigation.

AR/VR Systems

- Main Components.
  - Developer Tools.
  - Software Development Kit (SDK).
  - Device Hardware.
- Rendering Performance Counters.
  - Track the performance of AR/VR applications.
  - Normal user permission.
  - Hundreds of counters are available.

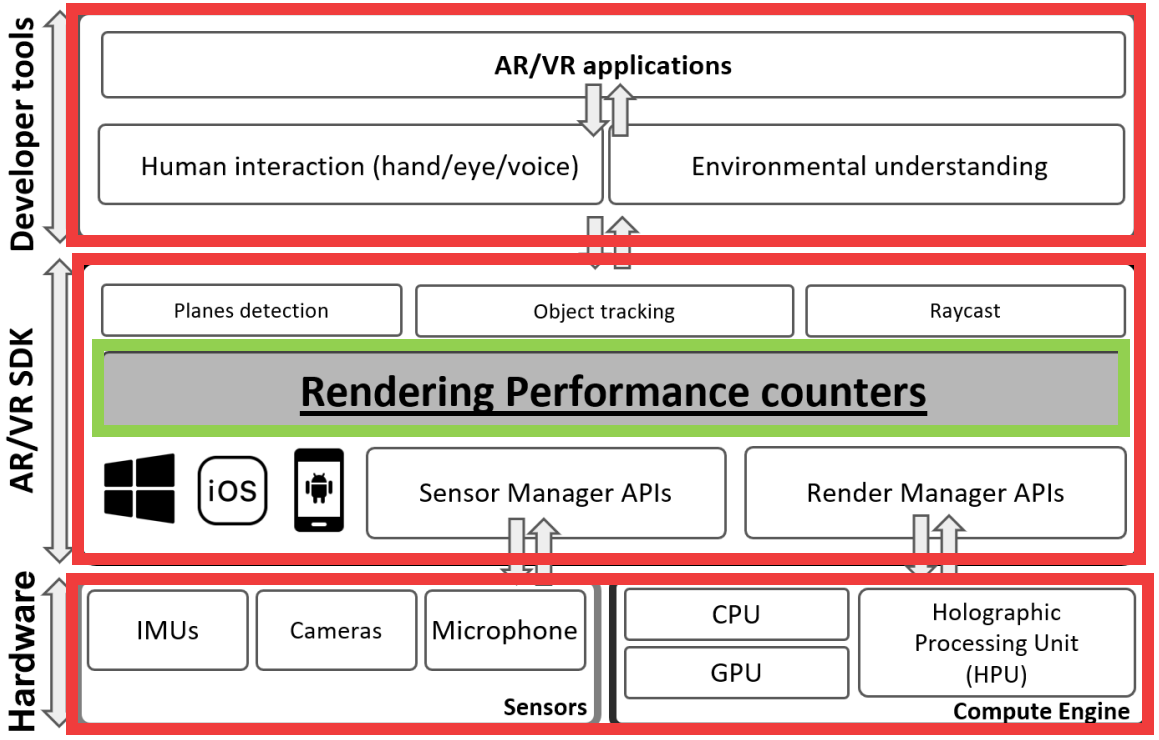

Threat model – Software Side-channel Attacks

- A malicious program runs in the background.
  - Standard application-level permissions.
  - No physical access.
  - Periodically probes rendering performance counters.

Leakage Vectors

- Memory Allocation API.
  - Expose memory usages on AR/VR devices.
- Rendering Performance Counters.
  - Unity & Unreal Engine SDK.
  - Frame Rate.
    - CPU/GPU frame rate, Refresh Rate, etc.
  - Thread Counters.
    - Game/Render thread time, etc.
  - Render Task Counters.
    - Number of draw calls, Number of primitives, Vertex count, etc.

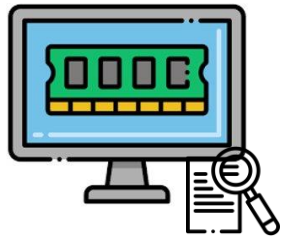

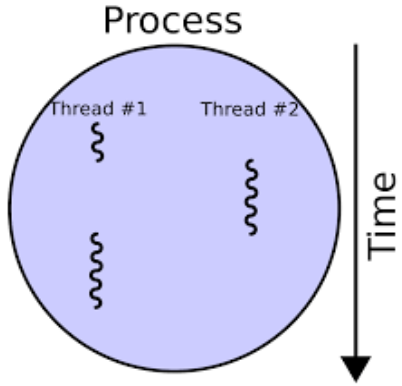

Side-channel Attacks Overview

- We demonstrate three classes of attacks.

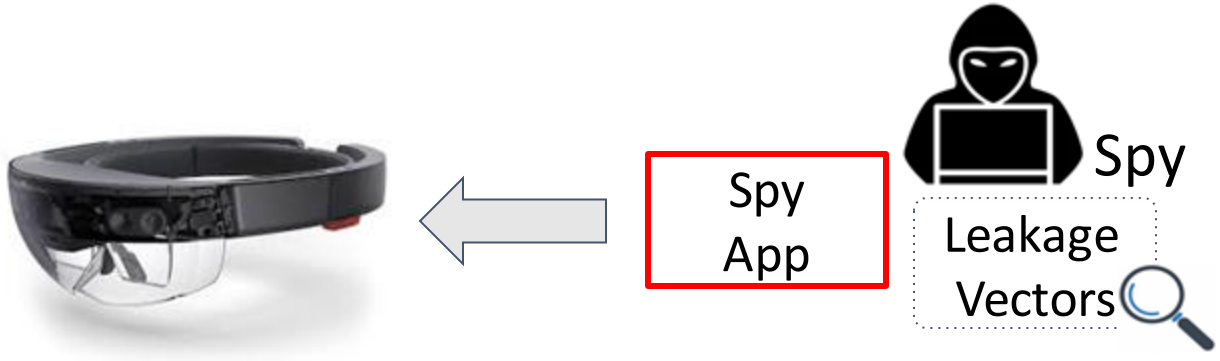

Spy application runs in the background.

Side-channel Attacks Overview

- We demonstrate three classes of attacks.
  - Spying on user interactions.

Spy application runs in the background.

Side-channel Attacks Overview

- We demonstrate three classes of attacks.
  - Spying on user interactions.
  - Spying on concurrent applications.

Concurrent applications

Spy application runs in the background.

Side-channel Attacks Overview

- We demonstrate three classes of attacks.
  - Spying on user interactions.
  - Spying on concurrent applications.
  - Spying on the real-world environment.

Real-world environment

Spy application runs in the background.

Experimental Setup

- Two popular headsets.
  - Microsoft Hololens 2 (AR).
  - Meta Quest 2 (VR).
- Ten volunteers.
  - Various ages, heights, weights, and gender.
- Attack workflow.

Microsoft HoloLens

oculus

Side-channel signal

Feature Engineering

Classifiers/regressors

Attack 1: Hand gestures inference

- **Victim:** Directly interacts with digital artifacts via hand gestures.
- **Spy:** Collects special signal patterns depending on victim's hand gestures.

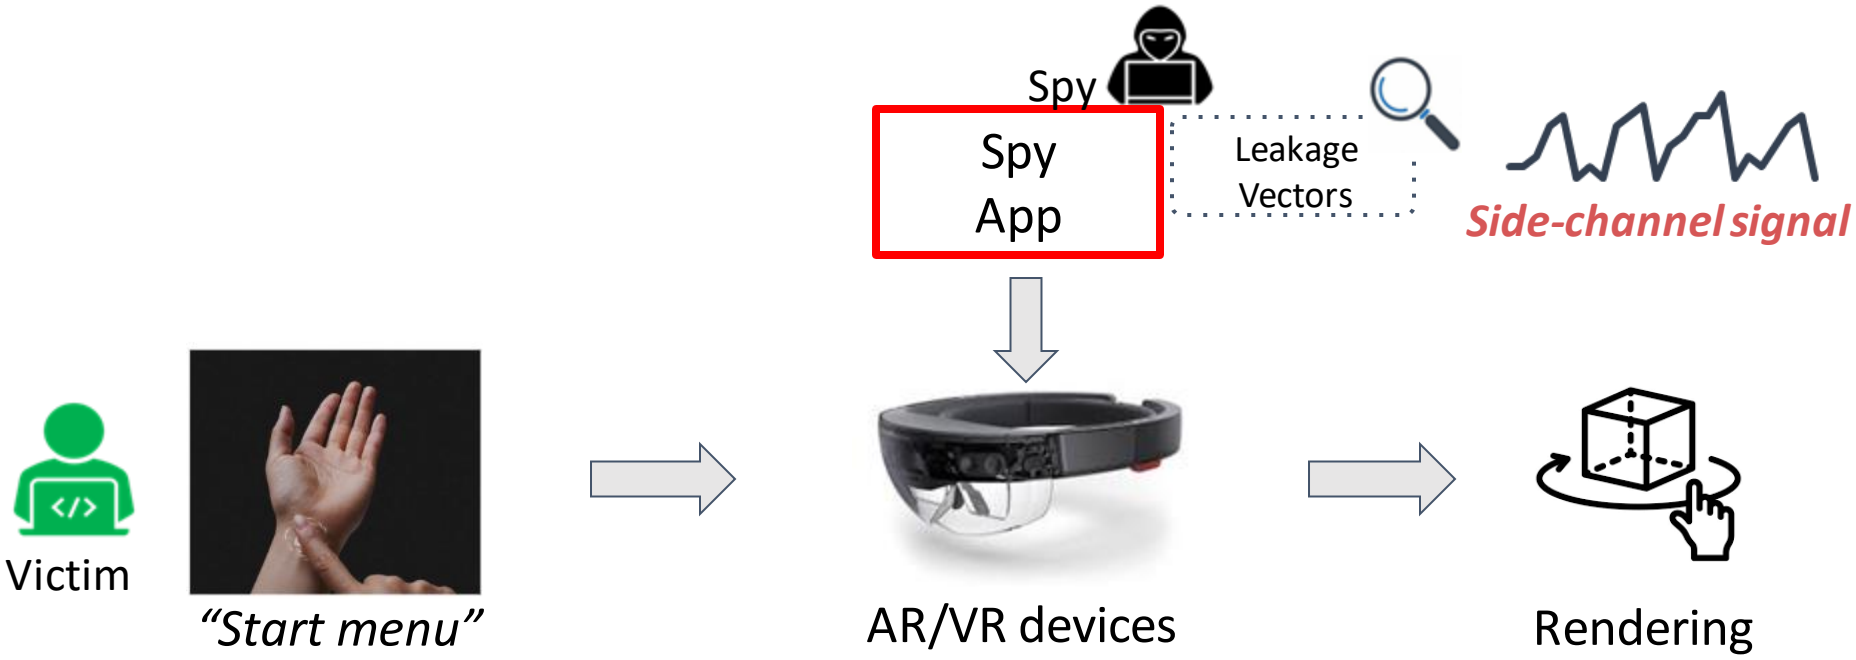

Performance Counter Trace

- *"Vertex count"*.
  - The number of vertices in existing 2D/3D scenes.

"Start Menu"

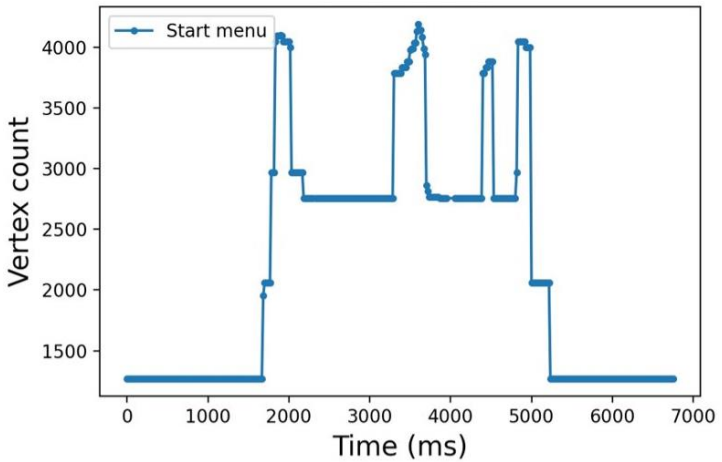

"Touch"

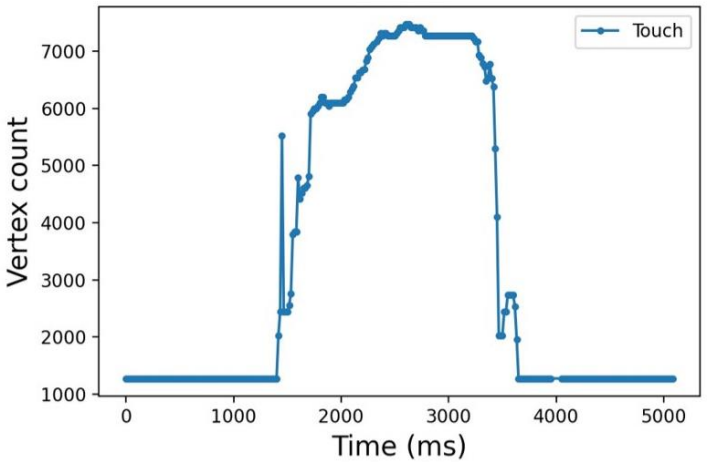

"Air Tap"

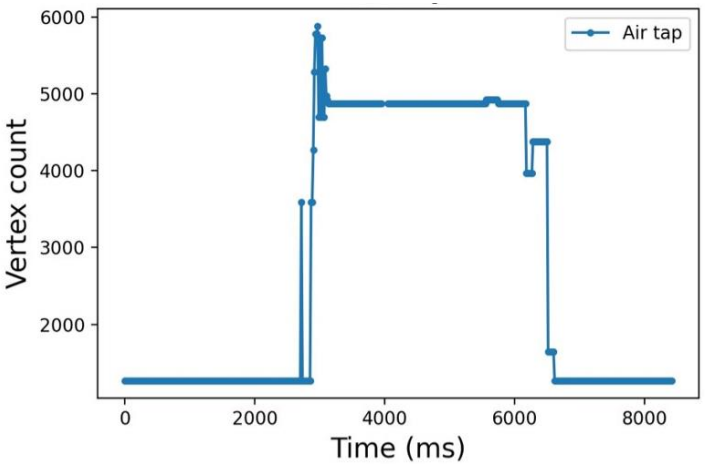

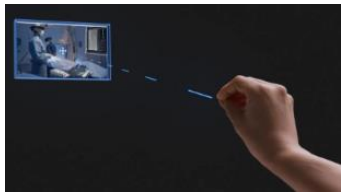

Classification Results

- Five basic system-level hand gestures on both Hololens 2 and Quest 2.
- The classification results for hand gestures inference attack:
  - K Nearest Neighbors (KNN).
  - Decision Tree (DT).
  - Random Forest (RF).
  - Light Gradient Boosting Machine (LightGBM).
  - Weighted majority rule voting (Voting).

	Hololens 2			Quest 2		
	F1	**Prec**	**Rec**	**F1**	**Prec**	**Rec**
KNN	53.6	55.4	54.2	57.9	58.3	58.8
DT	80.0	80.5	80.0	91.3	91.7	91.3
RF	86.6	86.6	86.7	**93.7**	93.8	93.7
LightGBM	84.7	86.7	85.0	89.0	91.9	90.0
Voting	**89.2**	89.3	89.2	91.3	91.9	91.3

Attack 2: Voice Commands inference

- **Victim:** Communicates with the headset through voice commands.

Attack 2: Voice Commands inference

- **Victim:** Communicates with the headset through voice commands.
- **Spy:** Measures a content-related pattern performed by the victim.

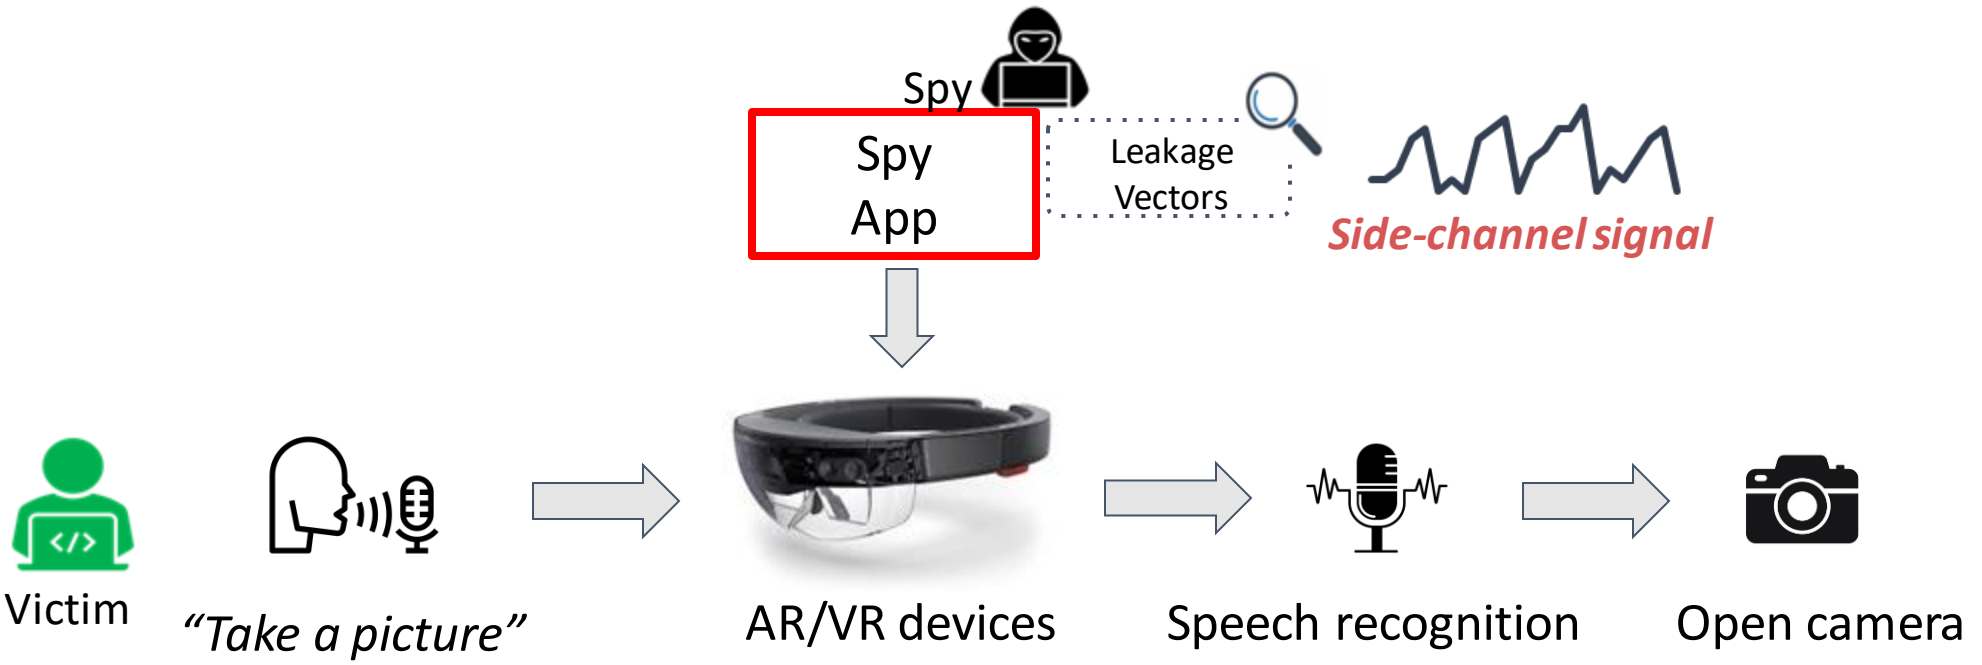

Memory Allocation Trace

- *"AppMemoryUsage" API*.
  - Track the spy app's current memory usage.

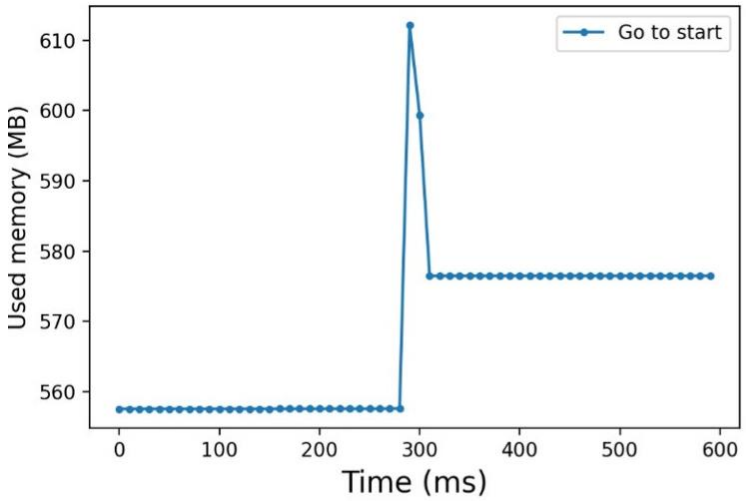

"Go to start"

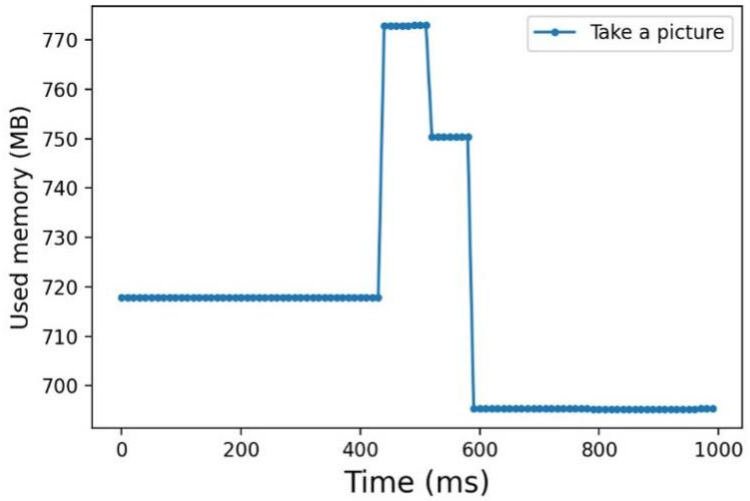

"Take a picture"

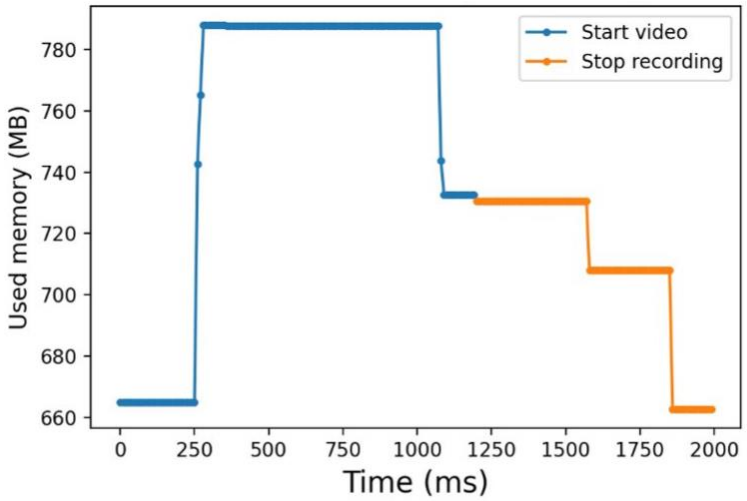

"Start/stop recording"

Classification Results

- Five basic headset-specific voice commands on both Hololens 2 and Quest 2.
- The classification results for voice commands inference attack:
  - K Nearest Neighbors (KNN).
  - Decision Tree (DT).
  - Random Forest (RF).
  - Light Gradient Boosting Machine (LightGBM).
  - Weighted majority rule voting (Voting).

	Hololens 2			Quest 2		
	F1	**Prec**	**Rec**	**F1**	**Prec**	**Rec**
KNN	87.5	87.7	87.5	65.9	73.3	62.0
DT	**93.7**	93.8	93.8	88.1	89.6	88.0
RF	91.2	91.3	91.2	86.0	89.3	86.0
LightGBM	88.9	90.9	89.5	90.3	93.0	90.8
Voting	91.3	92.4	91.3	**93.9**	94.0	94.0

Attack 3: Keystroke Monitoring

- **Victim:** Enters keystrokes through virtual keyboard.

"Foreground application"

"Virtual keyboard"

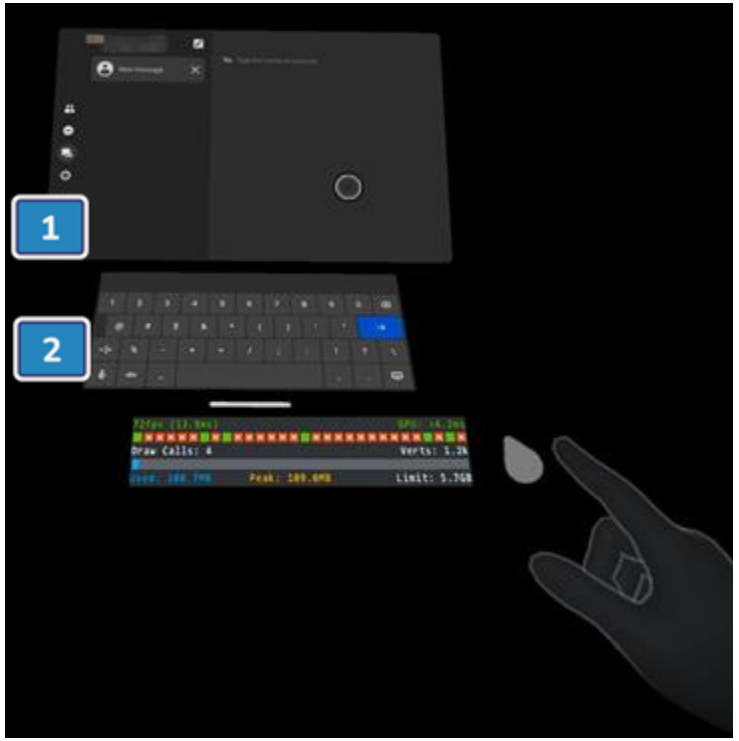

Attack 3: Keystroke Monitoring

- **Victim:** Enters keystrokes through virtual keyboard.
- **Spy:** Monitors rendering performance counters to infer the digit input of a victim.

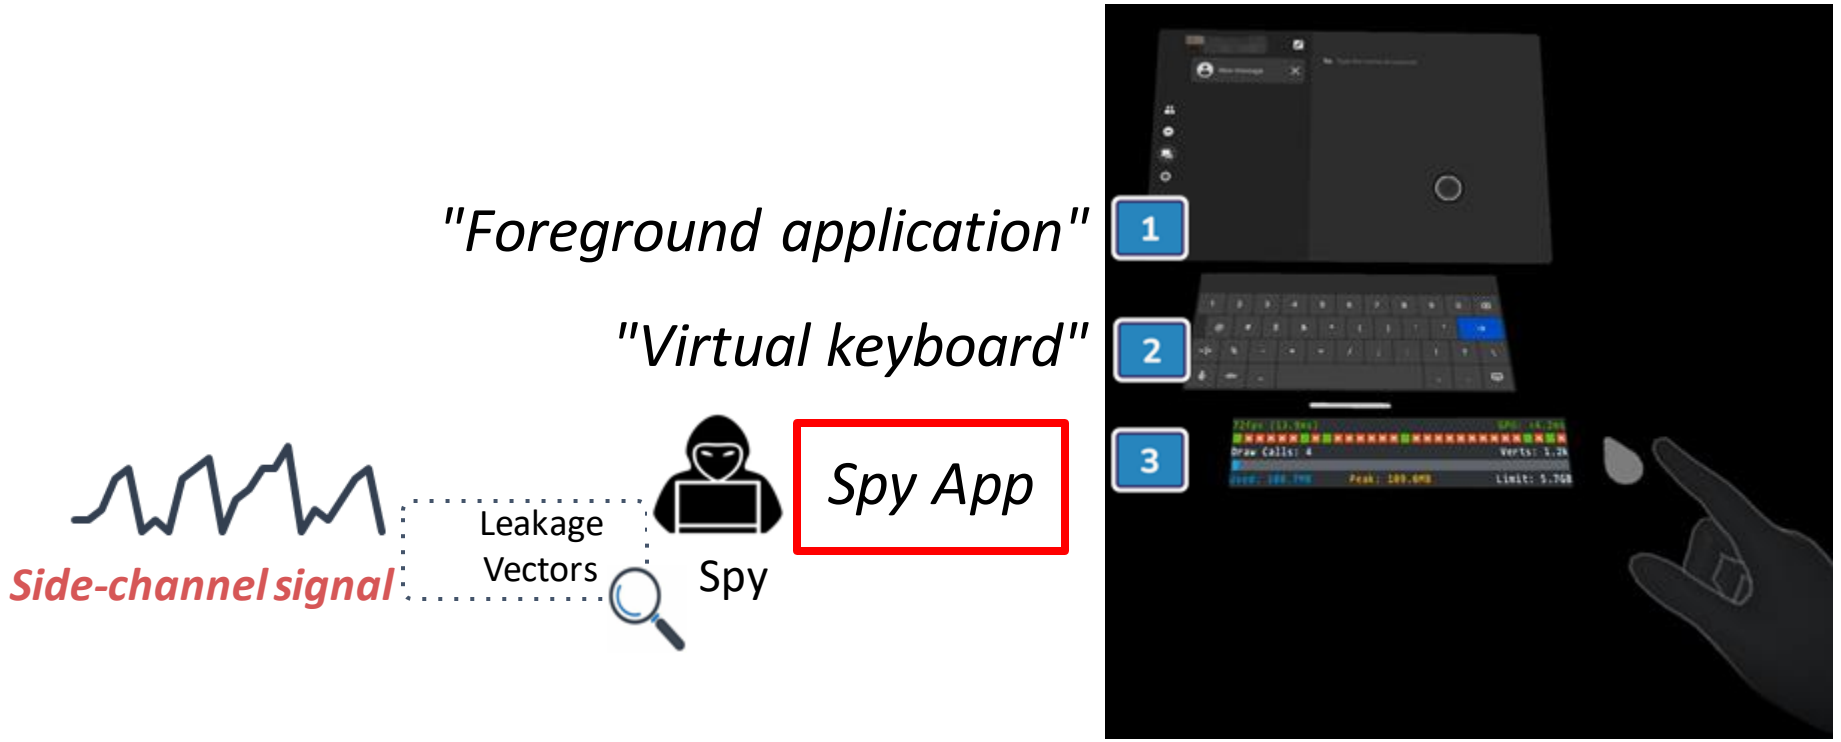

Performance Counter Trace

- *"Game thread time" & "Render thread time".*
  - Track the execution time of two primary threads in applications.

The victim presses the digit.

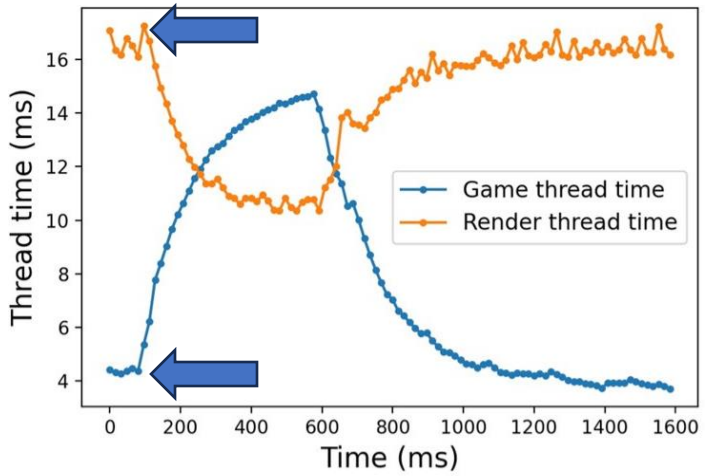

"Digit 0"

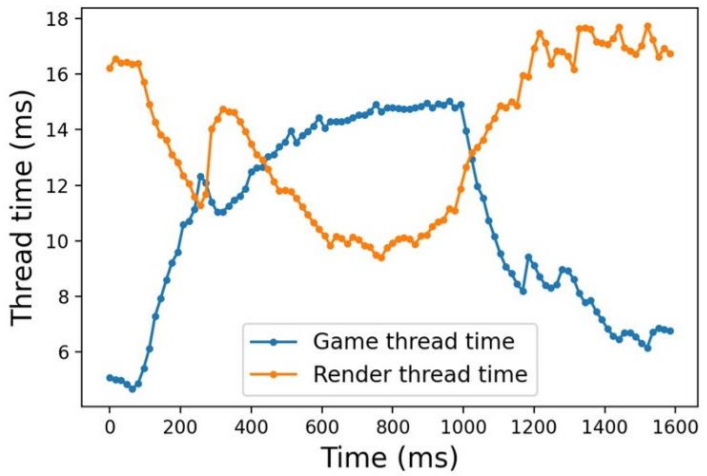

"Digit 2"

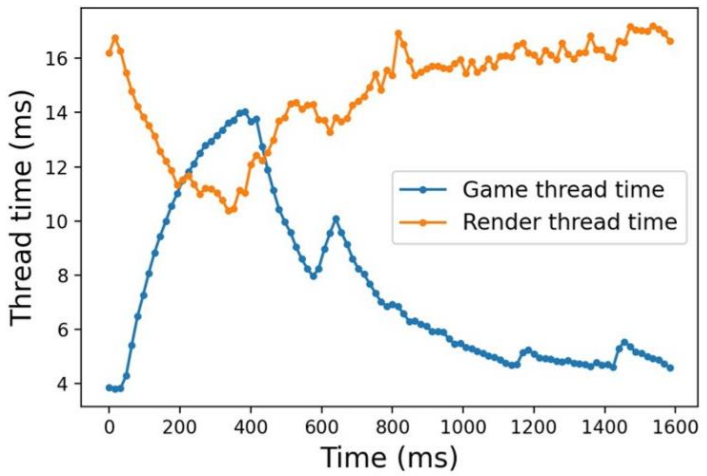

"Digit 9"

Classification Results

- Ten digits (0-9) on the virtual keyboard on both Hololens 2 and Quest 2.
- The classification results for keystroke monitoring:
  - K Nearest Neighbors (KNN).
  - Decision Tree (DT).
  - Random Forest (RF).
  - Light Gradient Boosting Machine (LightGBM).
  - Weighted majority rule voting (Voting).

	Hololens 2			Quest 2		
	F1	**Prec**	**Rec**	**F1**	**Prec**	**Rec**
KNN	43.3	49.6	44.3	44.1	49.4	44.0
DT	88.7	89.8	88.6	92.1	93.7	92.0
RF	52.1	54.0	52.9	73.7	75.5	75.0
LightGBM	87.5	88.0	88.8	**93.8**	94.8	94.0
Voting	**91.4**	91.7	91.4	90.1	91.6	90.0

Attack 4: Concurrent App Fingerprinting

- **Victim:** Launches a concurrent App on AR/VR devices.
- **Spy:** Track performance counters and identify the victim's application.

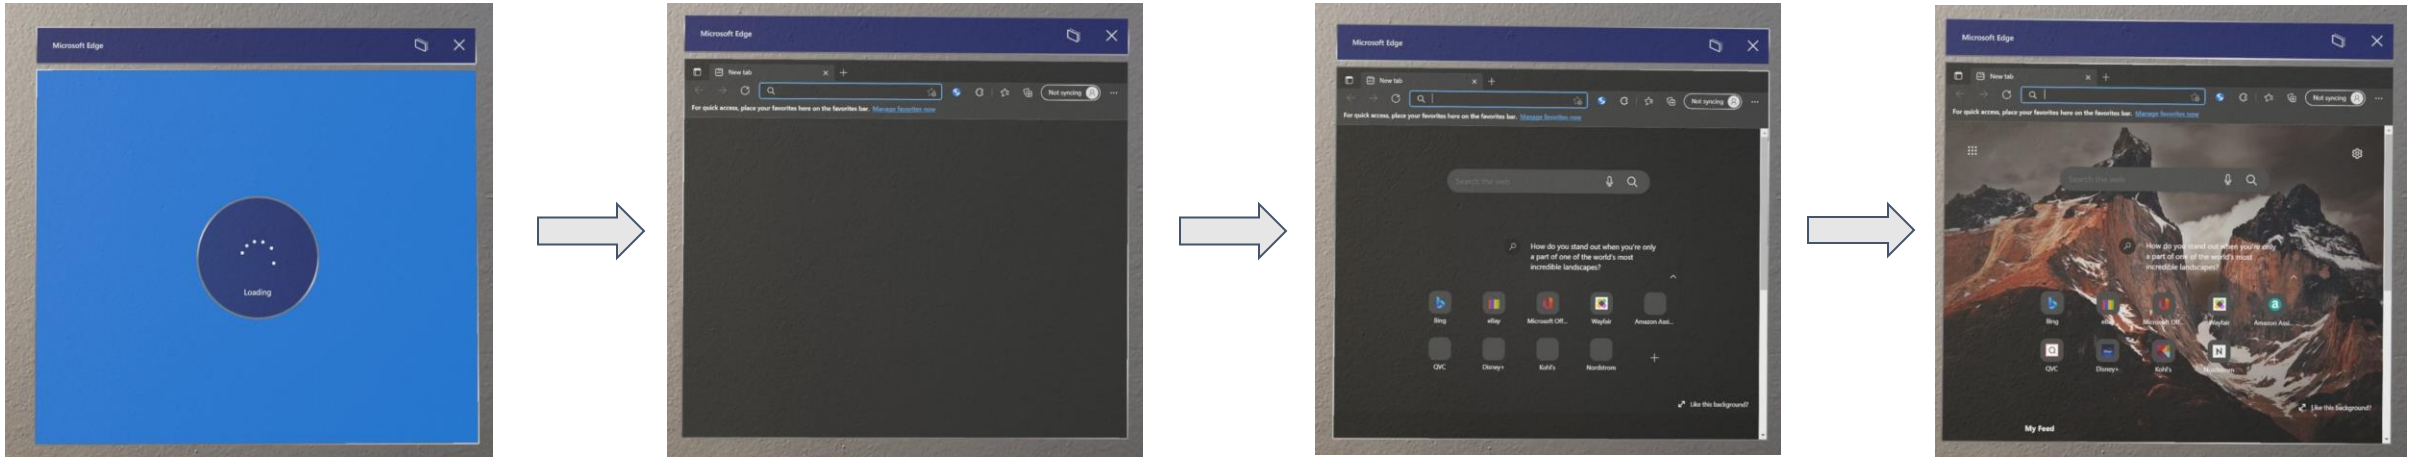

Time to launch an App on AR/VR devices

Performance Counter Trace

- *"Frame time"*.
  - Time takes for two consecutive frames are shown.

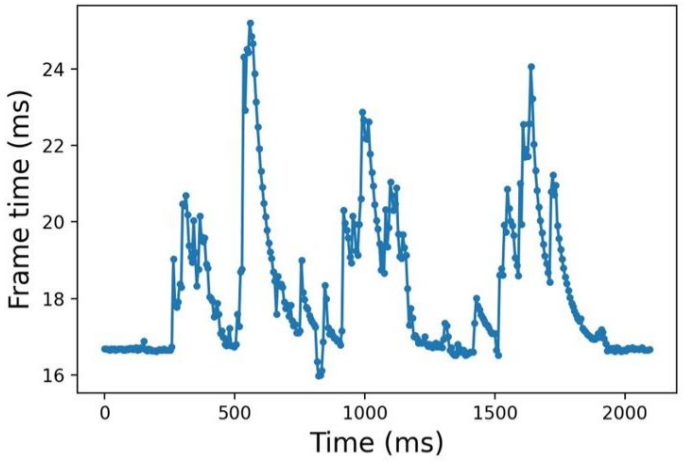

"Microsoft Edge"

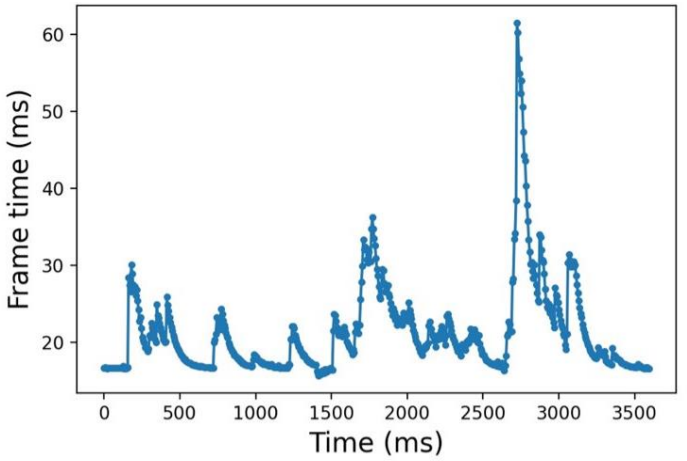

"One Drive"

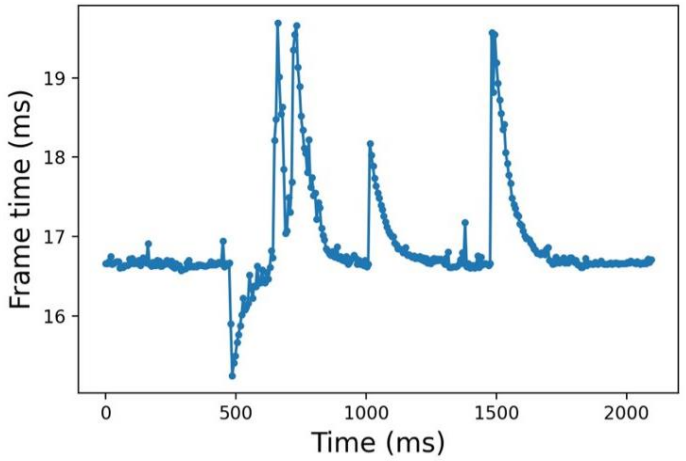

"Mail"

Classification Results

- Twelve applications are profiled on Hololens 2.
- The classification results for concurrent App fingerprinting:
  - K Nearest Neighbors (KNN).
  - Decision Tree (DT).
  - Random Forest (RF).
  - Light Gradient Boosting Machine (LightGBM).
  - Weighted majority rule voting (Voting).

	F1	Prec	Rec
KNN	33.7	39.4	35.0
DT	84.7	86.5	85.0
RF	51.3	53.0	50.8
LightGBM	85.8	87.4	86.8
Voting	**89.3**	91.0	89.2

Attack 5: Bystander Ranging

- **Victim:** Bystander steps into the field of view of an AR/VR device.

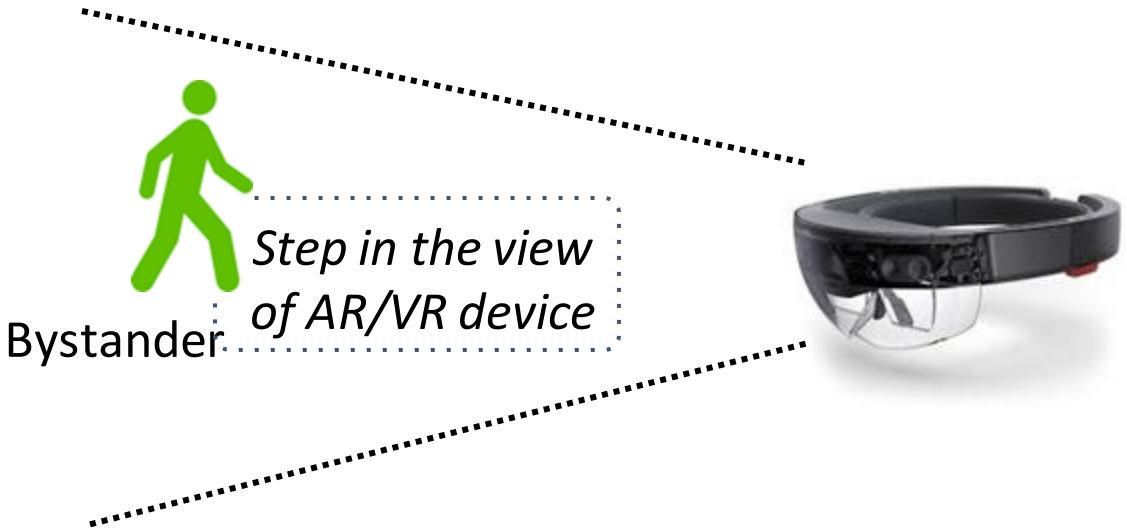

Attack 5: Bystander Ranging

- **Victim:** Bystander steps into the field of view of an AR/VR device.
- **Spy:**
  - Profiles leakage vectors.
  - Generates spatial mesh of the surrounding environment.

Turn an AR/VR device into a surveillance device!

Performance Counter Trace

- *"CPU frame rate"*.
  - CPU frame time between two consecutive frames.
- Distance-dependent fingerprint.

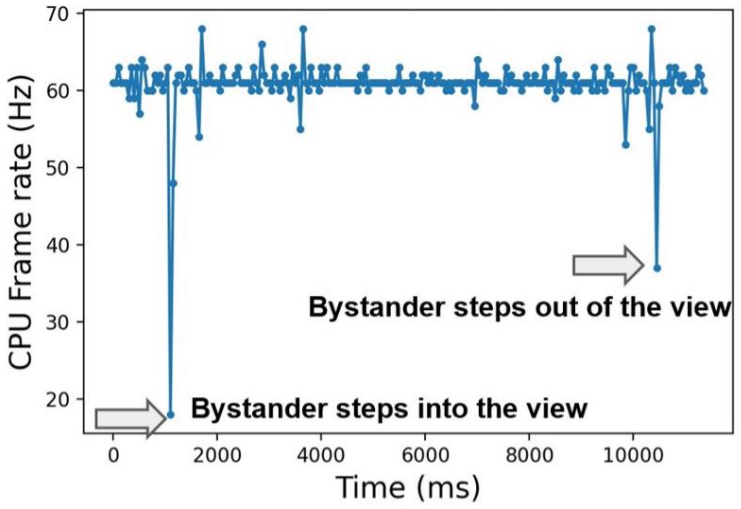

Distance = 0.5 meters

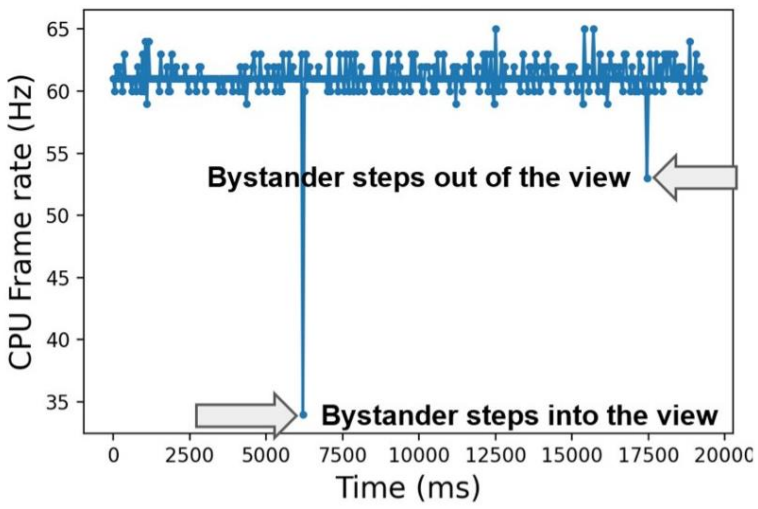

Distance = 2 meters

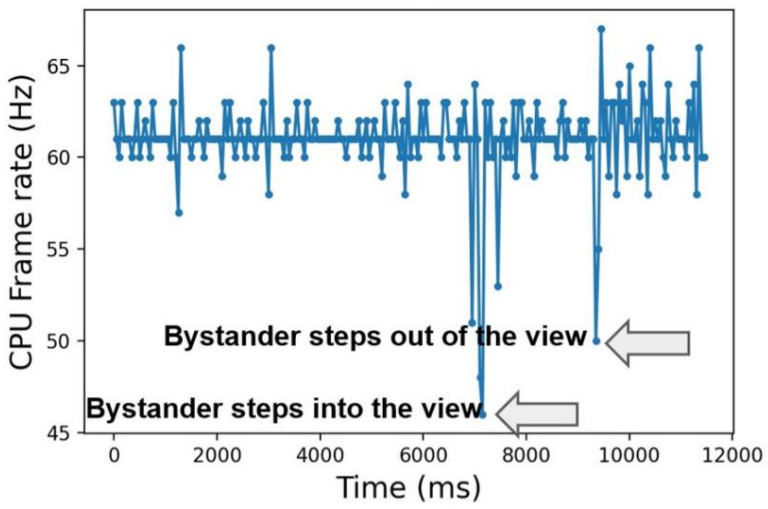

Distance = 4 meters

Regression Results

- Distance ranging from 0.5 meters to 5 meters (0.5, 1, 2, 3, 4, and 5 meters).
- The regression results (meters) for bystander estimation attack:
  - K Nearest Neighbors (KNN).
  - Decision Tree (DT).
  - Random Forest (RF).
  - Light Gradient Boosting Machine (LightGBM).
  - Weighted majority rule voting (Voting).

	KNN	DT	RF	LightGBM	Voting
MAE	0.401	**0.103**	0.257	0.279	0.164

10.3 cm

Mitigation

- Managing access to performance counters.
  - Completely blocking access to leaky APIs and counters.
  - Limiting the precision or rate of performance counters.

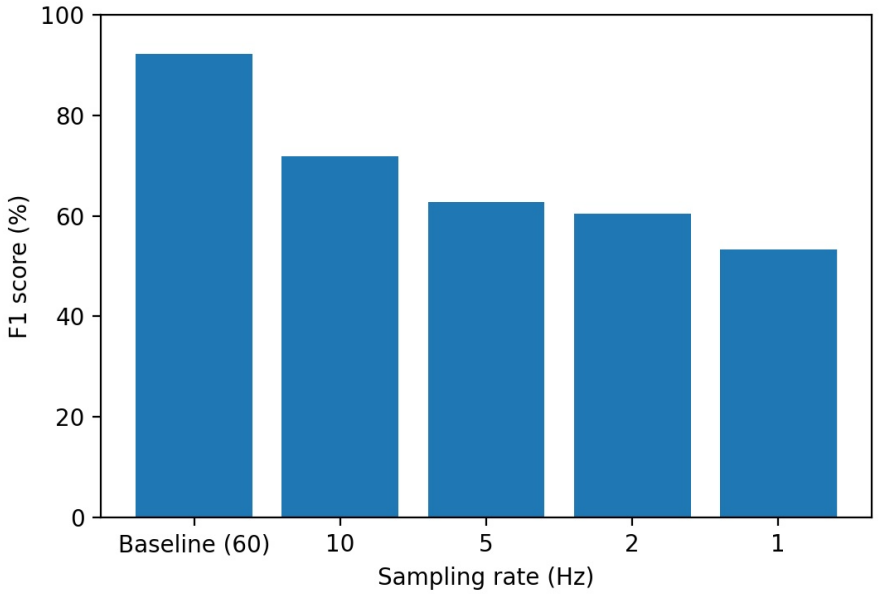

Conclusion

- Side-channels on AR/VR systems.
  - Through rendering performance counters (**First**).
- Three AR/VR-specific attack scenarios.
  - Five end-to-end side-channel attacks.
- Mitigation based on limiting the precision or rate is not effective.
- Future work:
  - Multi-user AR/VR systems; better profiling systems for AR/VR.

Thank you!

Any questions?

Yicheng Zhang

yzhan846@ucr.edu

https://yichez.site