Turning GPU into an FM Radio:
A Practical Data Exfiltration Framework from Air-gapped Systems

Rui Xiao[1,2], Sibo Feng[3], and Jinsong Han[3]

[1]Shanghai University of Finance and Economics, [2]MoE Key Laboratory of Interdisciplinary Research of Computation and Economics, [3]Zhejiang University

Presenter: **Dr. Yinghui He (NTU)**

INFOCOM 2026

Air-Gapped Systems

- *Physically isolated from external networks*
  - No wired or wireless network connection to the outside world
  - safeguarding sensitive data within the confines of the internal network

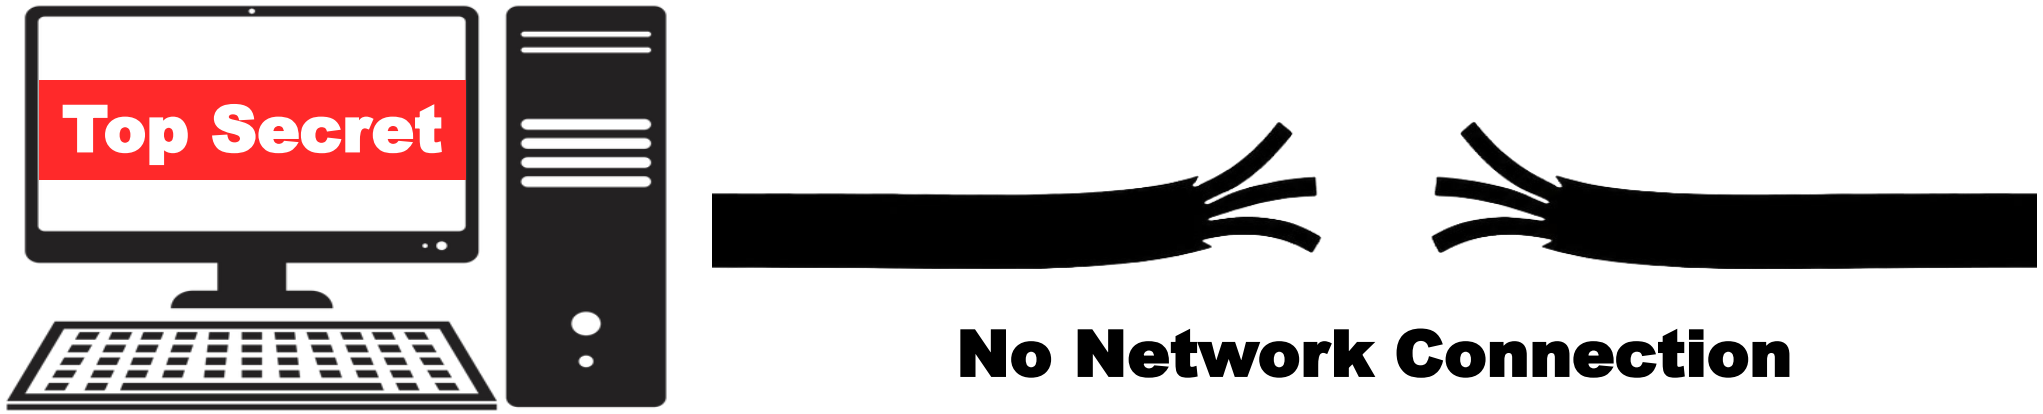

Air-Gapped Systems

- *Physically isolated from external networks*
  - No wired or wireless network connection to the outside world
  - safeguarding sensitive data within the confines of the internal network
  - Often deployed in ***highly sensitive*** environments

Google Distributed Cloud ***air-gapped***

NVIDIA NIM for Large Language Models: ***Air-Gap*** Deployment

"Many systems have been designed to operate in an ***air-gapped*** fashion"

Air-Gapped Systems

- *Physically isolated from external networks*
  - No wired or wireless network connection to the outside world
  - Often deployed in ***highly sensitive*** environments

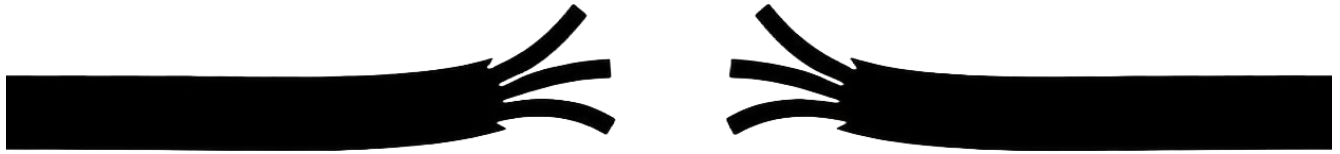

Question: Are Air-Gapped Systems really secure?

Prior work shows that even air-gapped systems **can be compromised** through **physical covert channels**

Physical Covert Channels for data exfiltration

- exploit *nontraditional artifacts* for communication
  - E.g., covert communication utilizing cooling fans

Static Fan: Bit 0

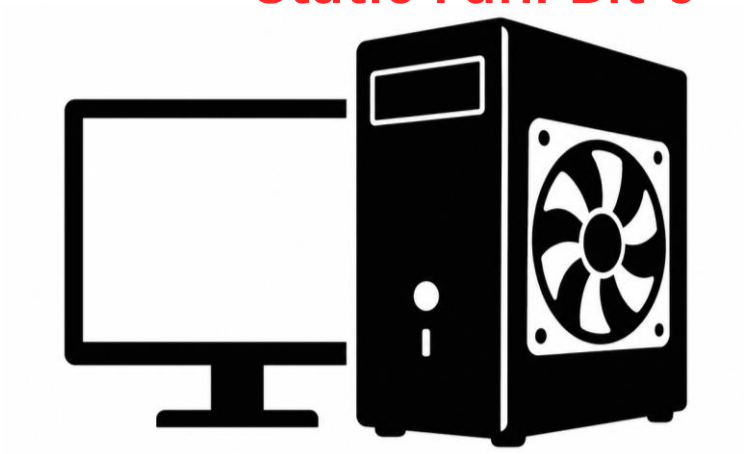

Fan Rotating Fast: Bit 1

Physical Covert Channels for data exfiltration

- exploit *nontraditional artifacts* for communication
  - E.g., covert communication utilizing cooling fans
- Other covert channels: thermal emanation, acoustics, electromagnetic side channels
- **Key Limitations**
  - low throughputs
  - line-of-sight (LoS) transmission
  - susceptibility to electromagnetic (EM) shielding
- render them **impractical in real-world scenarios**

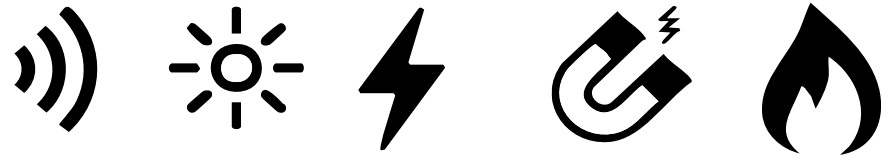

MagWhisper: A novel physical covert channel

- By modulating **GPU magnetic leakage**

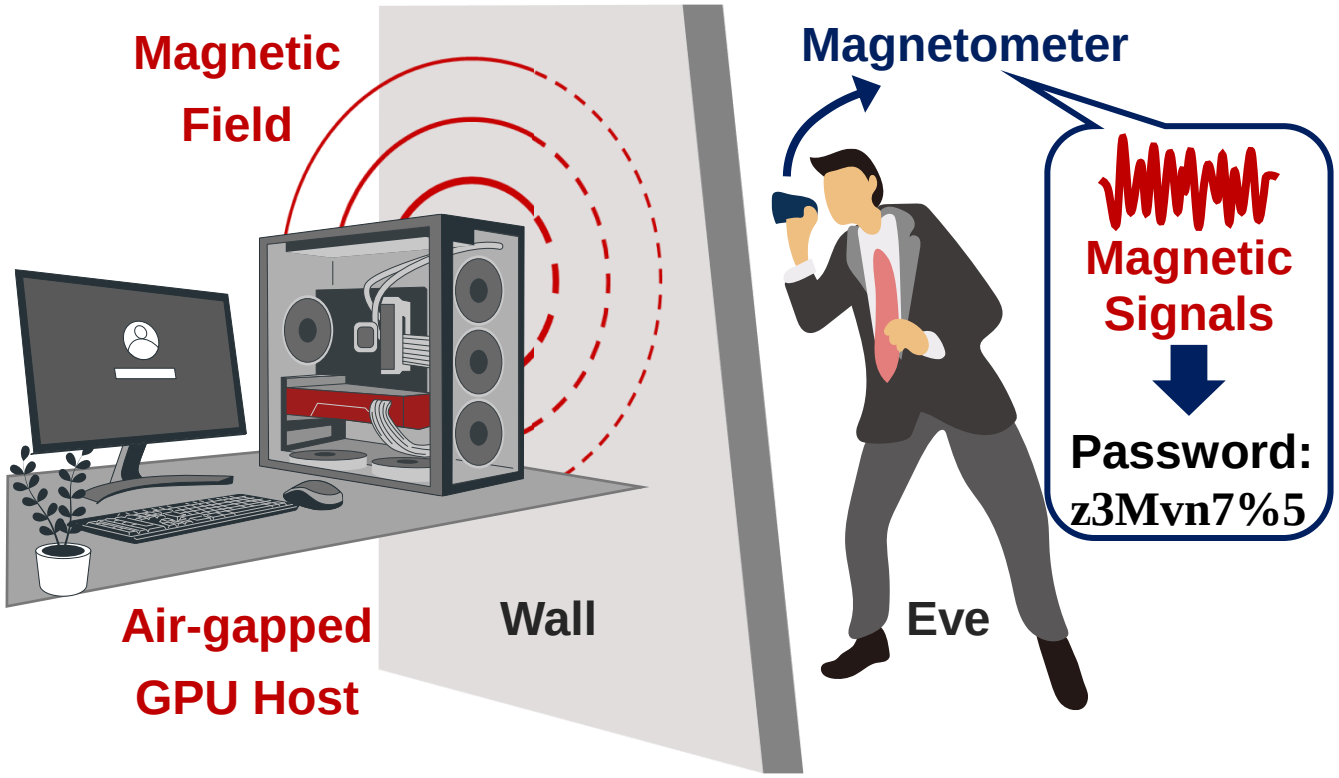

Key Features

- **Adequate Data Rate**
  - 133 bps
- **NLoS Transmission**
  - Cross-wall capability
- **Robust to EM Shielding**
  - Works in Faraday cage

GPU Magnetic Leakage Analysis

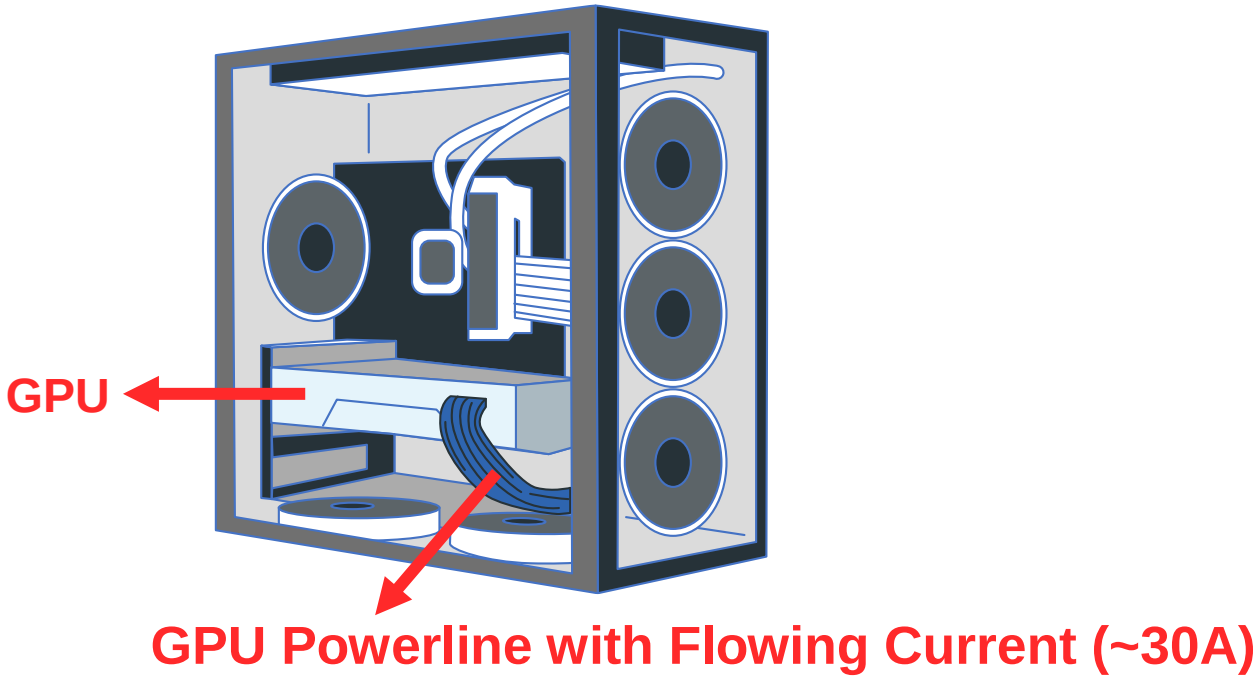

GPU Magnetic Leakage Analysis

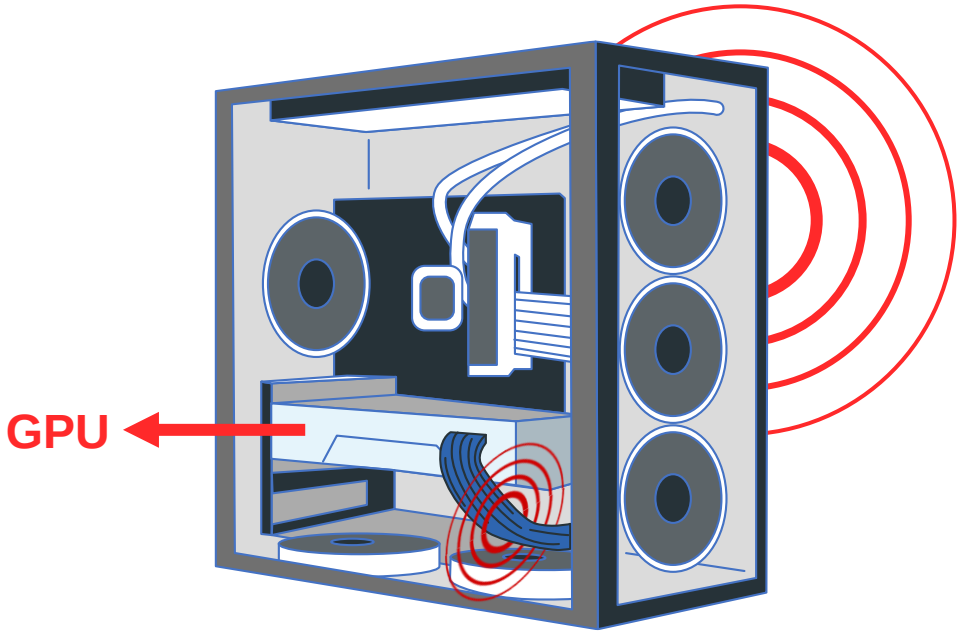

Attenuate with Distance

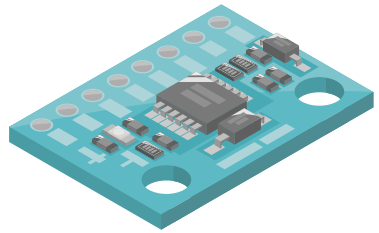

Remote Magnetometer
(Resolution ~1nT)

GPU Powerline with Flowing Current (~30A)

→ Induced *Magnetic Field* (~100μT)

GPU Magnetic Leakage Analysis

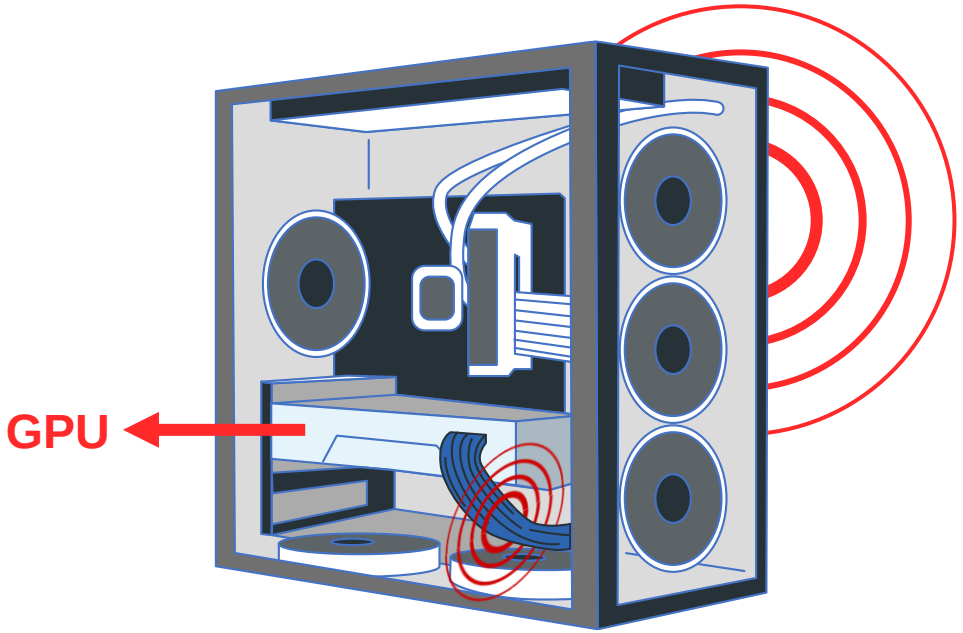

Attenuate with Distance

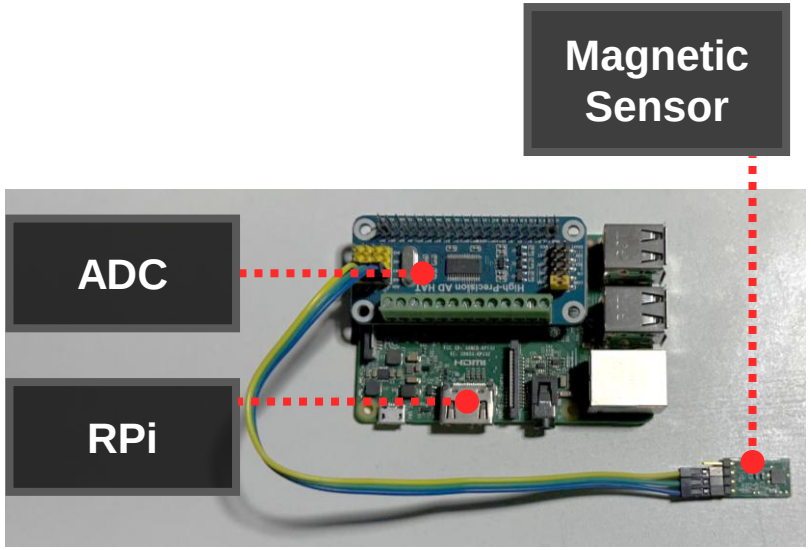

Receiver Prototype

GPU Powerline with Flowing Current (~30A)

→ Induced *Magnetic Field* (~100μT)

Challenge 1: Noisy Magnetic Channel

- **Ambient interference** with amplitude shifts and pulse-like noises

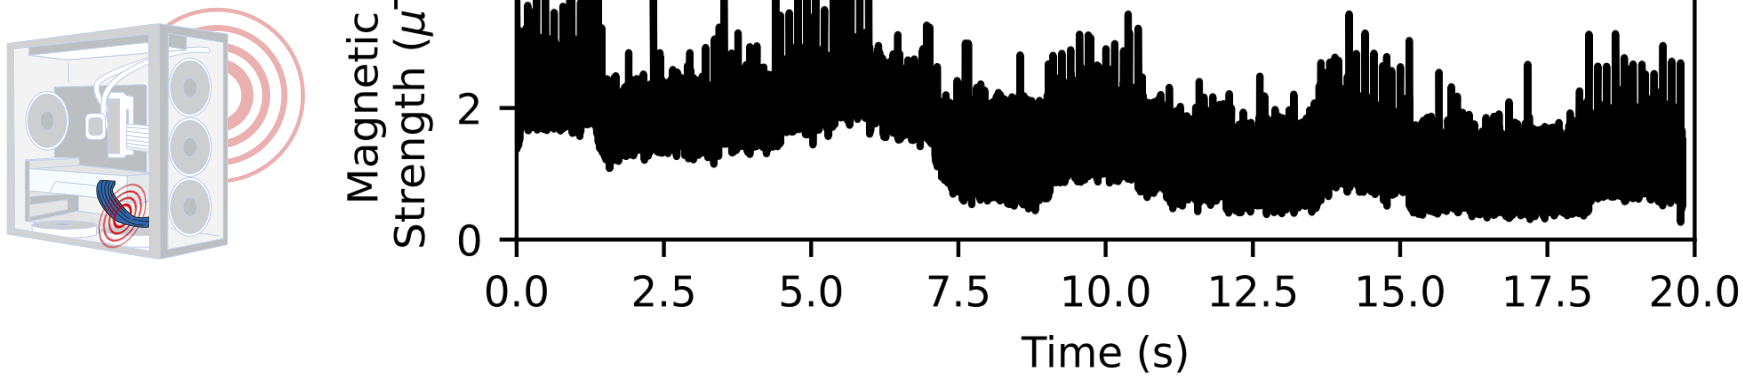

- Amplitude Modulation is thus shaky
- Use **Frequency-shift keying (FSK) modulation** instead
  - However, **switching between idle and active functions (Naïve FSK)** only produces low-frequency signals, due to the limited function-level switching rate
  - significantly restricts the data rate to **below 10 bps**

Solution: Efficient FSK with Operation-level Switching

- ***memory-dominant*** programs exhibit stable, periodic magnetic fluctuations

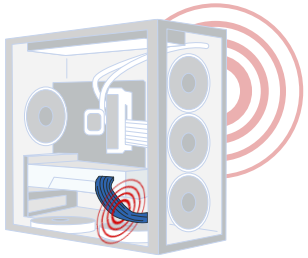

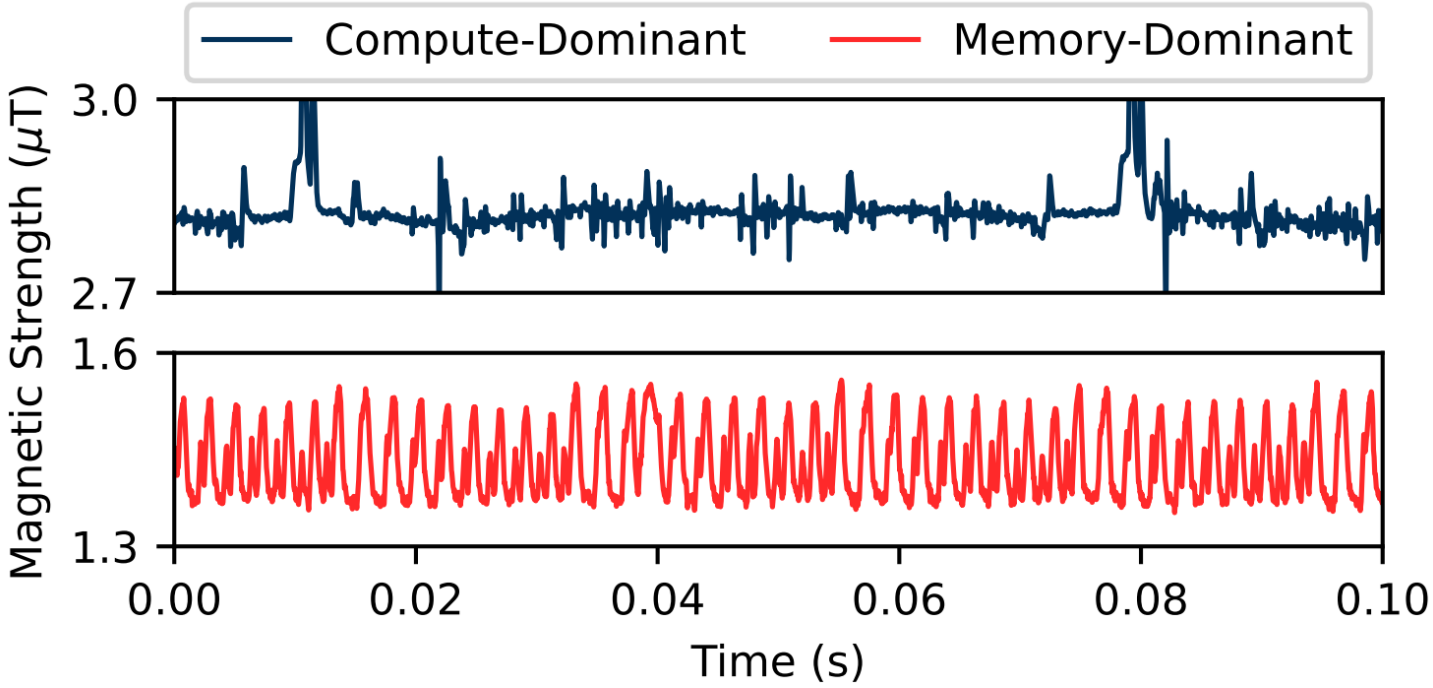

Solution: Efficient FSK with Operation-level Switching

- MagWhisper employs fine-grained **operation-level switching**, resulting in an adequate bit rate of 133 bps
- Magnetic **frequency can be controlled** by altering the number of GPU memory and computation operations

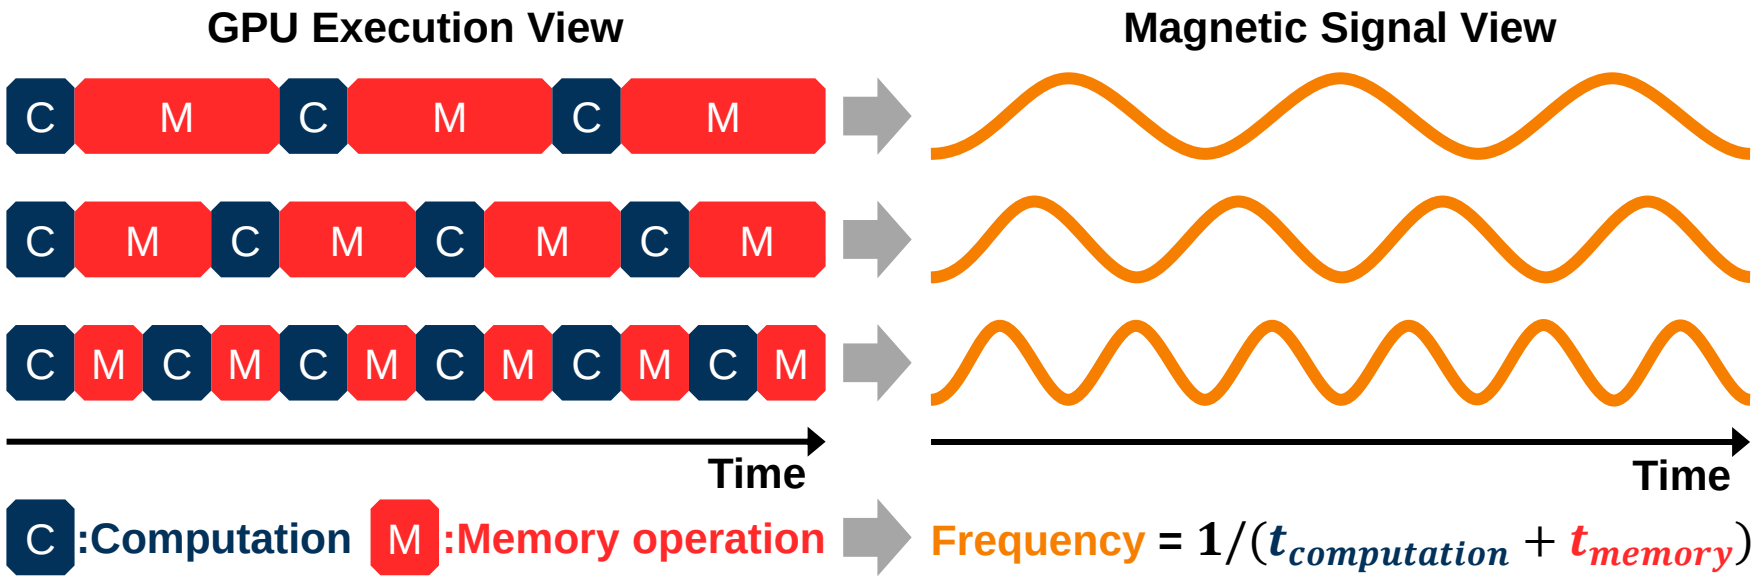

Challenge 2: GPU-dependent Magnetic Emanation

- Different Hardware Specifications of GPUs

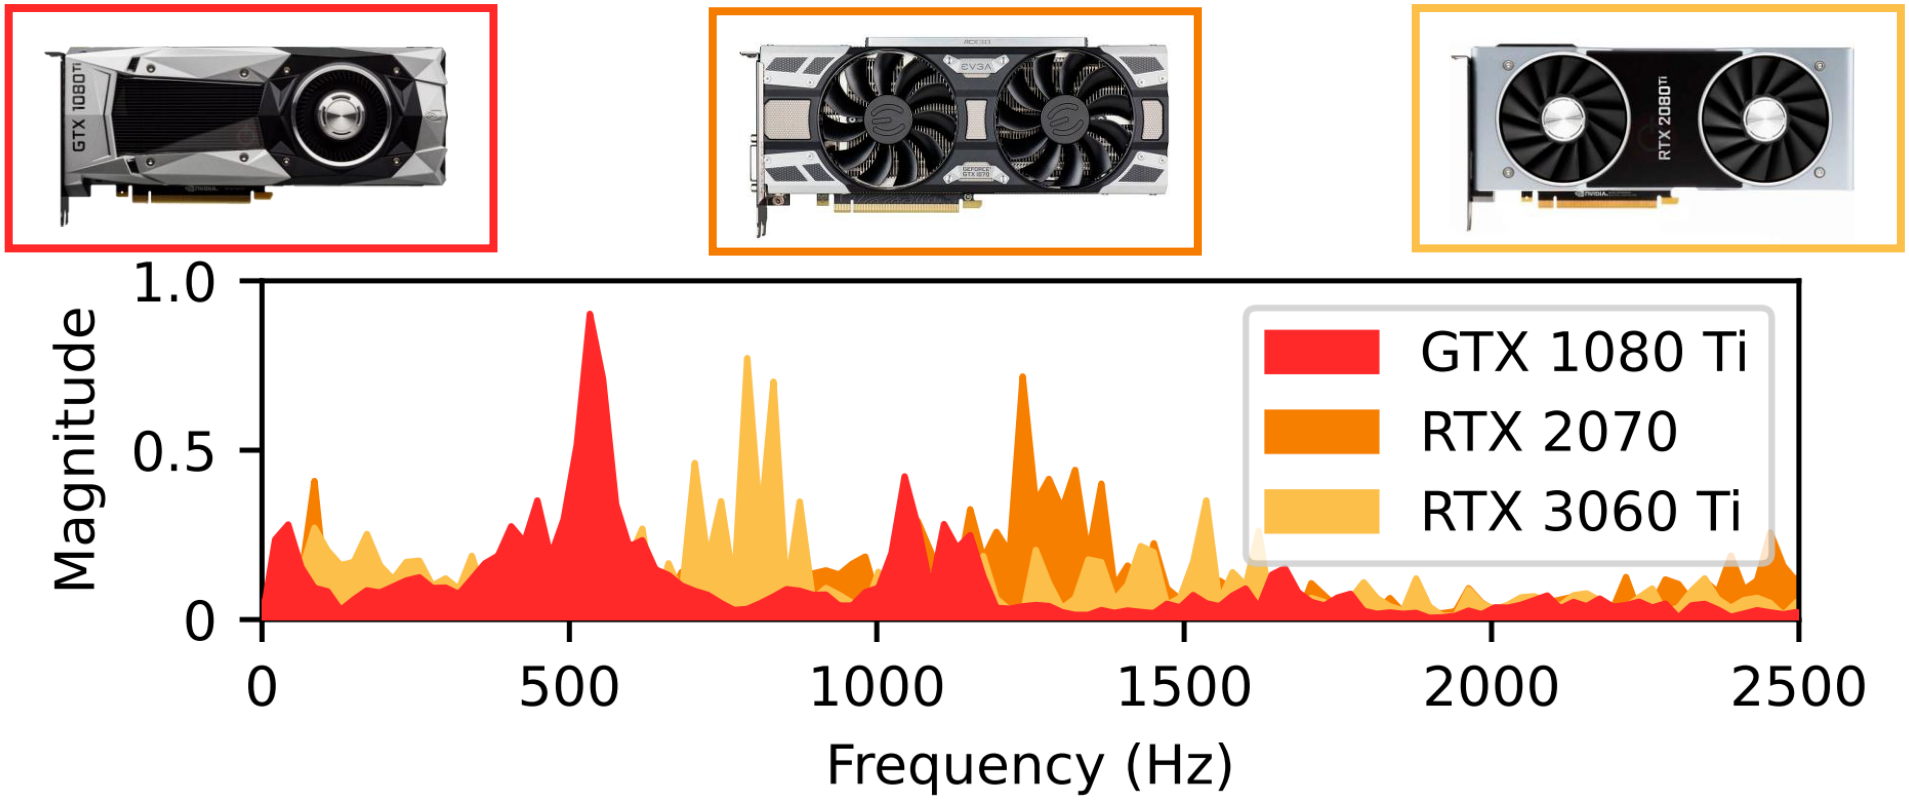

How to achieve robust and consistent demodulation

Solution: Channel Estimation with Preamble

- We draw inspiration from RF-based communication systems
- Packet structure in *MagWhisper*

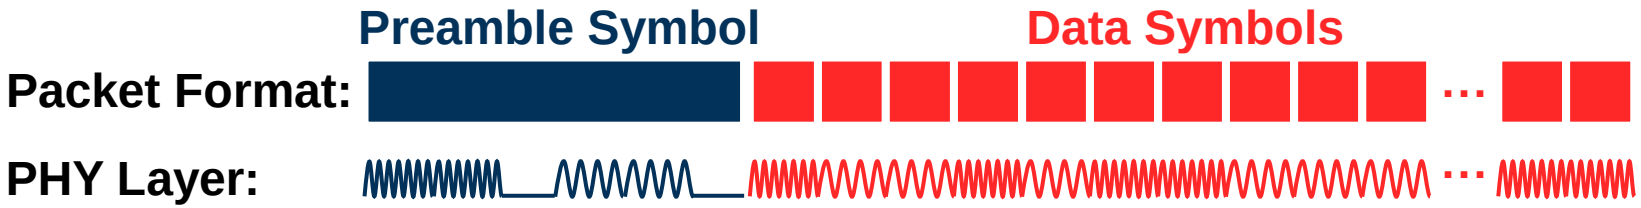

- *MagWhisper* can adaptively demodulated the data with detected preamble

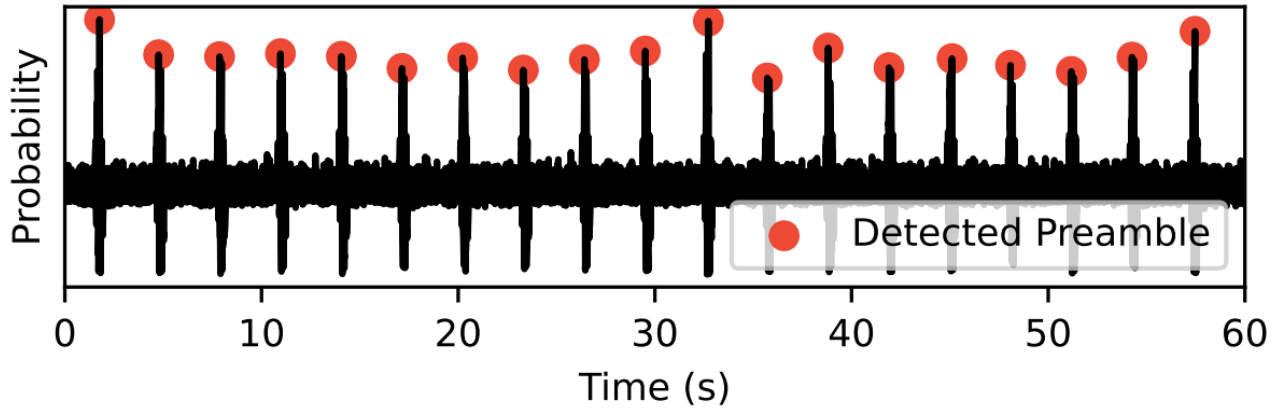

System Overview

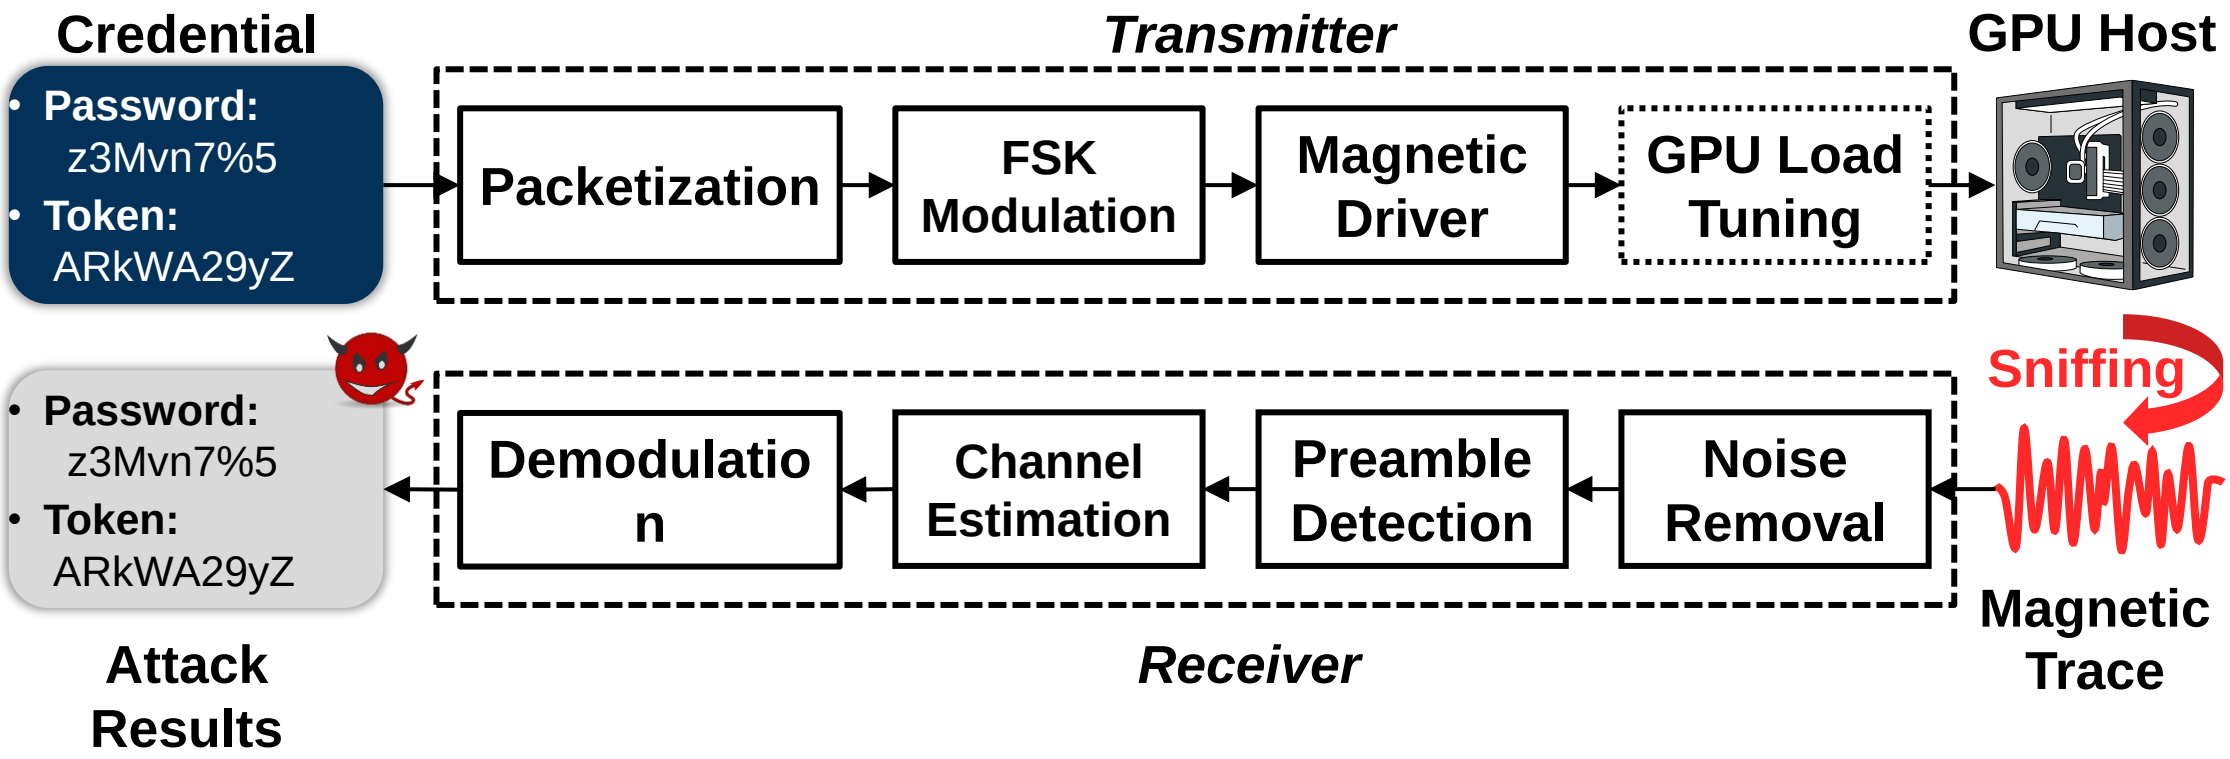

Evaluation

- **3** Heterogeneous GPUs (NVIDIA GTX 1080 Ti, RTX 2070, RTX 3060Ti)
- Overall performance
  - **bit rates** of 133 bps with a **bit error rate (BER)** lower than 1.2%

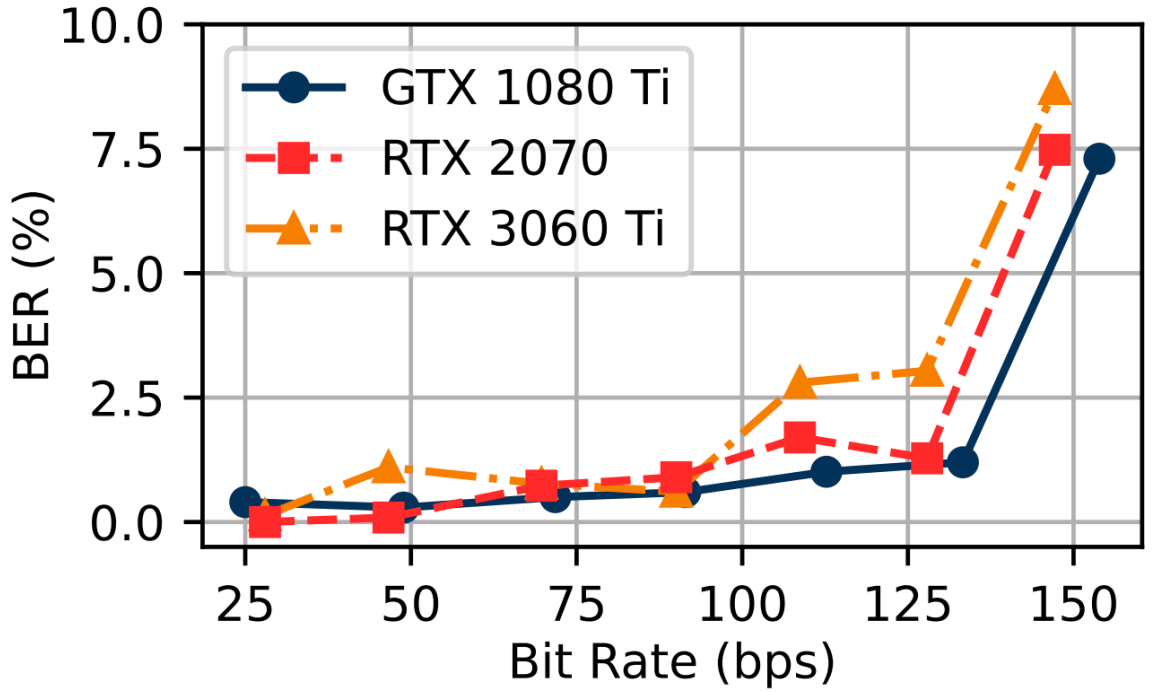

Evaluation

- **3** Heterogeneous GPUs (NVIDIA GTX 1080 Ti, RTX 2070, RTX 3060Ti)
- Overall performance
  - **bit rates** of 133 bps with a **bit error rate (BER)** lower than 1.2%
  - Works across different walls and at a **maximum distance of 35 cm**

(later we'll show the extension of 14 meters)

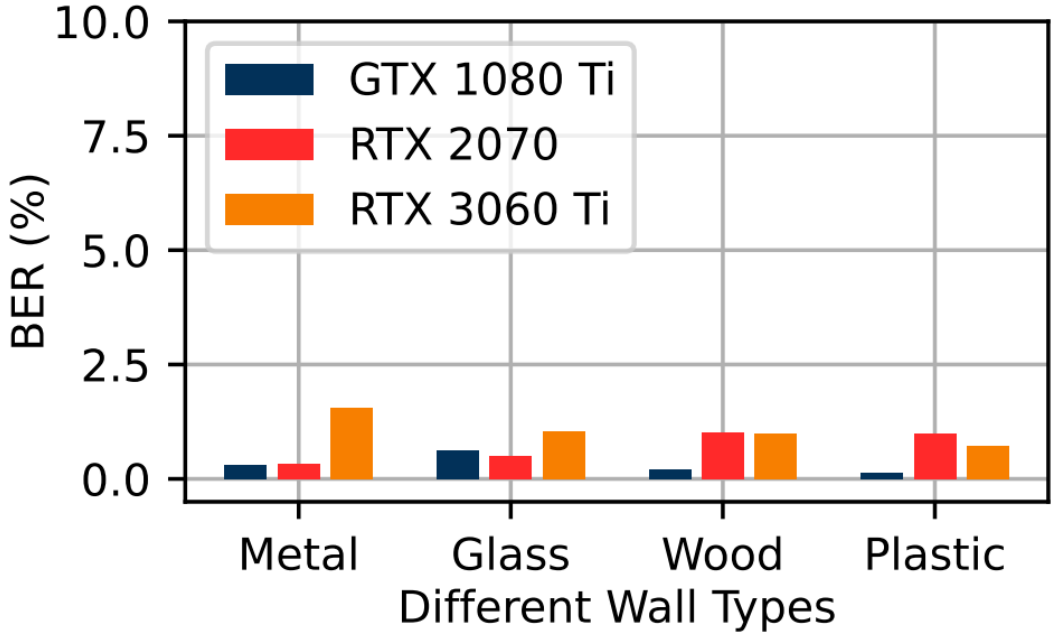

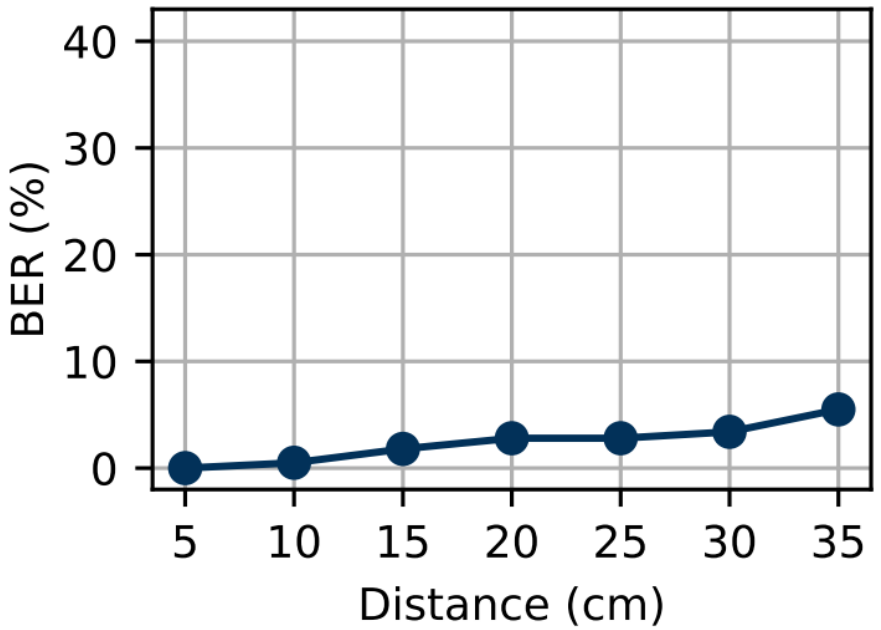

Evaluation

- **3** Heterogeneous GPUs (NVIDIA GTX 1080 Ti, RTX 2070, RTX 3060Ti)
- Overall performance
  - **bit rates** of 133 bps with a **bit error rate (BER)** lower than 1.2%
  - Works across different walls and at a **maximum distance of 35 cm**
  - **Case Study (1) Cross-Room** and **(2) Faraday Cage**

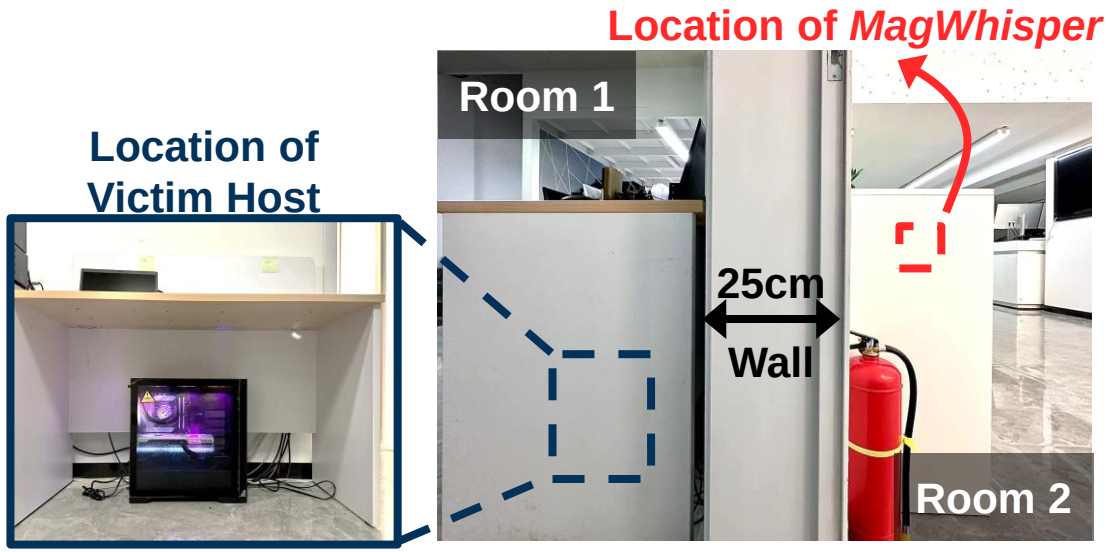

(a)

(b)

Evaluation

- **3** Heterogeneous GPUs (NVIDIA GTX 1080 Ti, RTX 2070, RTX 3060Ti)
- Overall performance
  - **bit rates** of 133 bps with a **bit error rate (BER)** lower than 1.2%
  - Works across different walls and at a **maximum distance of 35 cm**
  - **Case Study (1) Cross-Room** and **(2) Faraday Cage**

	Personal Information	Private Key
Ground Truth	Date of Birth: January 1, 1900 Address: 123 Main Street, Anytown	MIIFRzCCBC+gAwIBAg IQC2FFwm4d7zP4+7Z1 lACNrzANBgkqhkiG9w 0BAQsFADBH
Received (w/o FEC)	Date of Birth: JaNucr} 1, 1900 Address* 1"3 Main Stseet, Anytown	MIIFRzCCBC+gAwIBAG ISC2FFwm4d7zP4+7Z1 lASNrjANBgkqhkiF9w 0BAQsFADBH
Received (with FEC)	Date of Birth: January 1, 1900 Address: 123 Main Street, Anytown	MIIFRzCCBC+gAwIBAg IQC2FFwm4d7zP4+7Z1 lACNrzANBgkqhkiG9w 0BAQsFADBH

Faraday Cage
(Victim Host Inside)

***MagWhisper* in the drawer**

(b)

Discussion: Long-range extension of MagWhisper

- Same GPU transmitter, no modification
- GPU activity **also modulates GHz-range EM leakage**
- Captured using USRP B210 + directional antenna
  - **133 bps** transmission rate
  - < 0.7% BER **at 14 m**
  - Drawback: can be EM shielded

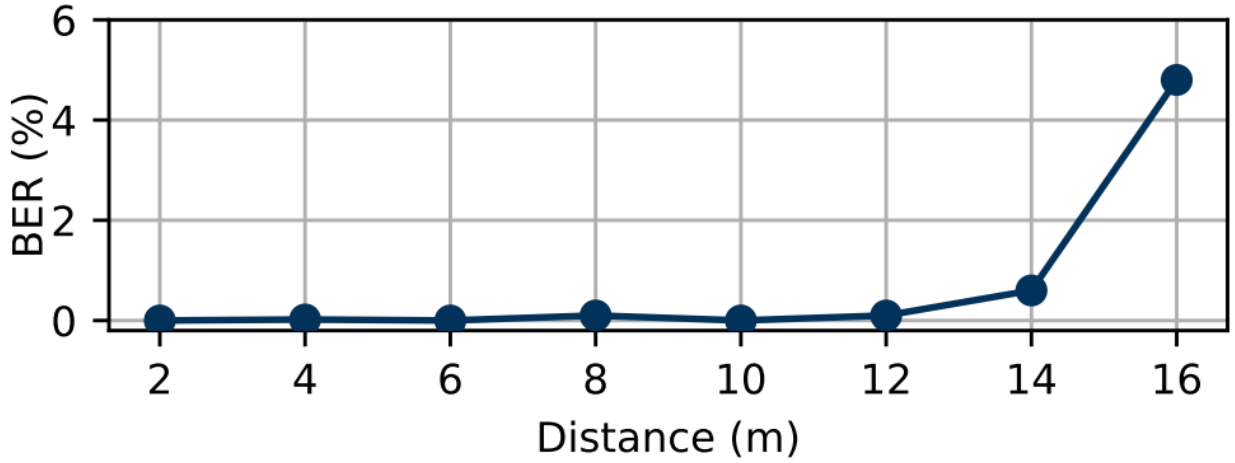

Conclusion

- We introduce ***Modelspy***, a long-range **DNN architecture snooping attack** utilizing the **GPU EM side channel**.
- We hope to highlight that securing advanced AI systems requires rethinking the **physical security of AI infrastructure**.

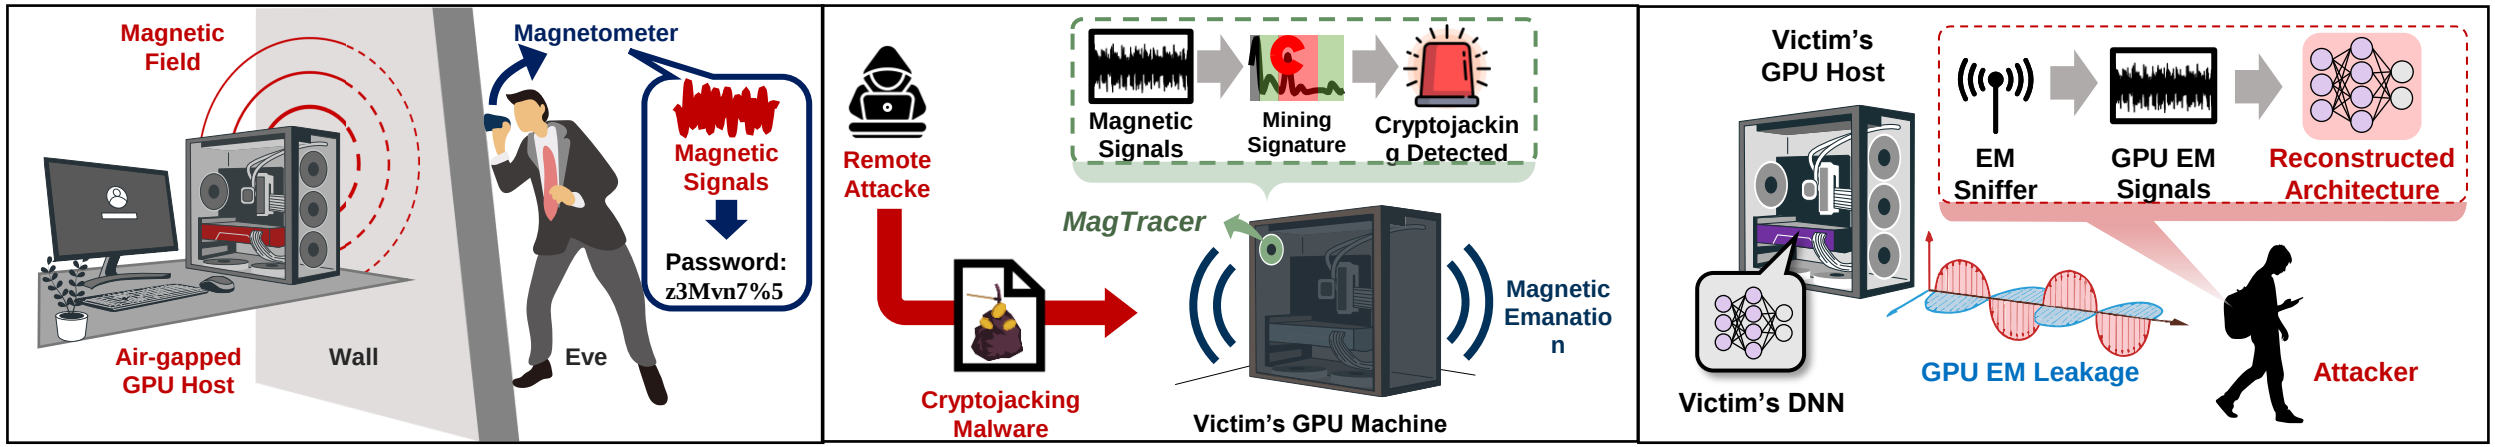

INFOCOM'26 (This Work) **MobiCom'23** **NDSS'26**

Interested in GPU Physical Side Channels?
I'd be happy to discuss.

Thank you!

Homepage

Rui Xiao

Assistant Professor at SUFE

Wireless, Mobile, Sensing, Security

Please feel free to contact me!

Mail: ruixiao24@sufe.edu.cn