PRATEEK SAHU

PhD Candidate, Computer Architecture
University of Texas at Austin

prateeks@utexas.edu
https://www.linkedin.com/in/sahuprateek/
https://prateeksahu.github.io

SUMMARY

Defensive security researcher specializing in cross-stack security research - from hardware microarchitectural side-channels and runtime ransomware detection to vulnerabilities in modern AI inference pipelines. Published at MICRO and ASPLOS, with ongoing work on AI threat red-teaming and adversarial attack amplification across the software-hardware boundary.

Research Interests: AI Security, Systems Security, Malware Detection, Microarchitectural Side Channels

EDUCATION

Ph.D., Computer Architecture ▪ *University of Texas at Austin* 2020 - Present (Expected June 2026)
Advisor: Prof. Mohit Tiwari
Dissertation: Identifying and Detecting Cross Layer Threats in Modern Workloads
Qualcom Innovation Fellowship recipient (2022)

M.Sc, Computer Architecture ▪ *University of Texas at Austin* 2017 - 2019
GPA: 3.79/4

Bachelor of Engineering, Electrical ▪ *Indian Institute of Technology, Kanpur* 2011 - 2015
GPA: 8.2/10
Kishore Vaigyanik Protsahan Yojana (KVPY) Fellowship, Dept. of Science and Technology, India

RESEARCH EXPERIENCE

Graduate Research Assistant ▪ *SPARK Lab, UT Austin* ▪ *Prof. Mohit Tiwari* Spring 2018 - Present

AI Systems Security Research (2023 - Present)

- AI Attack RedTeam – Cascade framework demonstrating how software-hardware attack gadgets compose to amplify adversarial threats in AI systems; exemplified using fault injection and model jailbreak [pdf]
- AI Threat Exploration – Systematization of hardware and software threat vectors in AI inference pipelines [pdf]
- Investigated prompt injection vulnerabilities in M365 production deployments to suppress correct responses and misattribution of ownership using malicious RAG injection [pdf]

Ransomware Detection Research (2021 - 2023)

- Designed a lifecycle-aware ransomware detector leveraging cross-layer signals – hardware performance counters, OS events, and network telemetry – for early-stage detection with reduced data loss
- Trained detection model on MITRE ATT&CK malware lifecycle stages; demonstrated higher detection confidence vs. single-layer approaches

Micro-architectural Side Channel Research (2018 - 2021)

- Developed speculative interference attacks that exploit younger speculativeinstructions affecting older instruction latency, undermining existing cache defense mechanisms [pdf]
- Created Cyclone, a side-channel malware detector that tracks cyclic contention events across security domains, evaluated against cache-based side and covert channels on the gem5 simulation environment [pdf]

WORK EXPERIENCE

Graduate Technical Intern ▪ *Intel Labs* Jun - Nov 2020, May - Aug 2021

- Built proof-of-concept for secure data offload pipeline in Intel SGX-adjacent accelerator environments.
- Designed remote attestation and key provisioning flows for confidetial compute offload use cases.

Engineering Intern ▪ *Qualcomm Technologies* May - Aug 2019

- Architectural design for SMT in QDSP6 Control Unit and their security implications.
- Formal verification of existing RTL design to find hardware scheduler bugs.

Member of Technical Staff ▪ *VMware India* July 2015 - June 2017

- Filed 3 patents with USPTO for a system health monitoring tool using vSphere metrics.
- Designed and implemented a private cloud resource and cost monitoring tool using utilization statistics.
- Developed containerized micro-services CI/CD pipelines for vSphere application infrastructure.

SELECTED PUBLICATIONS

- Cascade: Composing Software-Hardware Attack Gadgets for Adversarial Threat Amplification in Compound AI Systems (arXiv '26)
  Sahu, P.*, Banerjee, S.*, Vahldiek-Oberwagner, A., Sanchez, J.V., Tiwari, M. [pdf]
- SoK: A systems perspective on compound AI threats and countermeasures (arXiv '24)
  Sahu, P.*, Banerjee, S.*, Luo, M., Vahldiek-Oberwagner, A., Yadwadkar, N.J., Tiwari, M. [pdf]
- Understanding Sidecars in Cloud Orchestration (SESAME '25, co-located with EuroSys '25)
  Sahu, P., Wei, S., Yadwadkar, N.J., Tiwari, M. [pdf]
- Speculative Interference Attacks: Breaking Invisible Speculation Schemes (ASPLOS '21)
  Behnia, M., **Sahu, P.**, Paccagnella, R., Yu, J., Zhao, Z., Zou, X., Unterluggauer, T., Torrellas, J., Rozas, C., Morrison, A., Mckeen, F., Liu, F., Gabor, R., Fletcher, C.W., Basak, A., Alameldeen, A. [pdf]
- Cyclone: Detecting Contention-Based Cache Information Leaks Through Cyclic Interference (MICRO '19) Harris, A., Wei, S., **Sahu, P.**, Kumar, P., Austin, T., Tiwari, M. [pdf]

TECHNICAL SKILLS

AI/ML: PyTorch, LLM inference pipelines, RAG systems, agentic frameworks
Languages: C, C++, Python, Bash, x86 Assembly
Simulation: gem5, QEMU
System Tools: Perf, eBPF, Intel PT, Valgrind