The University of Texas at Austin
Electrical and Computer Engineering
Cockrell School of Engineering

Using **Power-Anomalies** to Counter **Evasive Micro-Architectural Attacks** in Embedded Systems

Shijia Wei,

Aydin Aysu*, Michael Orshansky, Andreas Gerstlauer, and Mohit Tiwari

ECE, North Carolina State University*

ECE, The University of Texas at Austin

New Threats in Embedded Systems

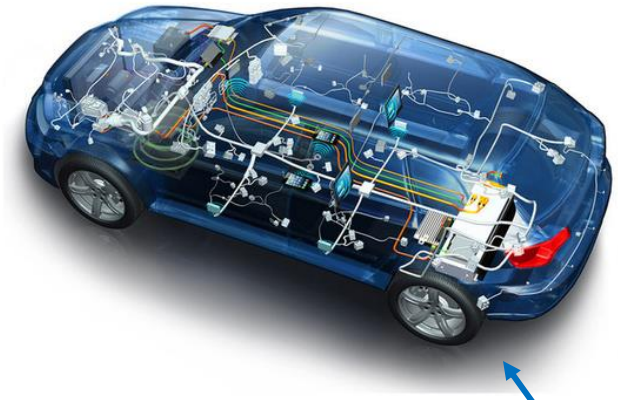

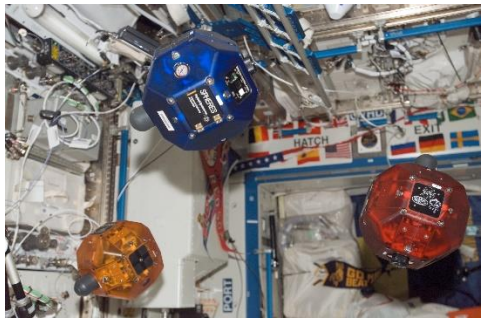

Long-lasting, hard to patch

Unforeseen attacks

Novel micro-arch. attacks

Bypass generic defenses

Covert Channel

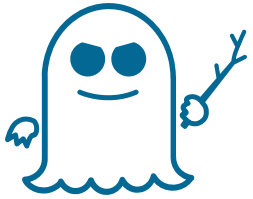

Spectre

Rowhammer

In-Band Solutions to Specific Attacks

- Custom defenses for
  - Covert-Channels
  - Spectre
  - Rowhammer

Application/Compiler	[Retpoline, Google] [Bitmask, WebKit] [FuzzyTimer, WebKit/Chromium/Mozilla]	
Operating Systems	[KAISER, ESSoS'17] [EPTI, ATC'18] [CATalyst, HPCA'16]	[CATT, USENIX'17] [ZebRAM, OSDI'18] [ANVIL, ASPLOS'16]
Processors/SoC	[PLCache&RPCache, ISCA'07] [CEASER, MICRO'18] [DAWG, MICRO'18] [Microcode, Intel]	[ARMOR, Patented] [Kim et al., ISCA'14]

Out-of-Band Detection via Power: Intuition

- Using the power side channel for defense
  - Coarse-grained external power measurement
  - Generic, upgradable, tamper-proof, against unknown attacks
- Workloads have unique power signature
  - Regular behavior in embedded applications
  - Clearly differentiable behavior of attacks

Face Detection

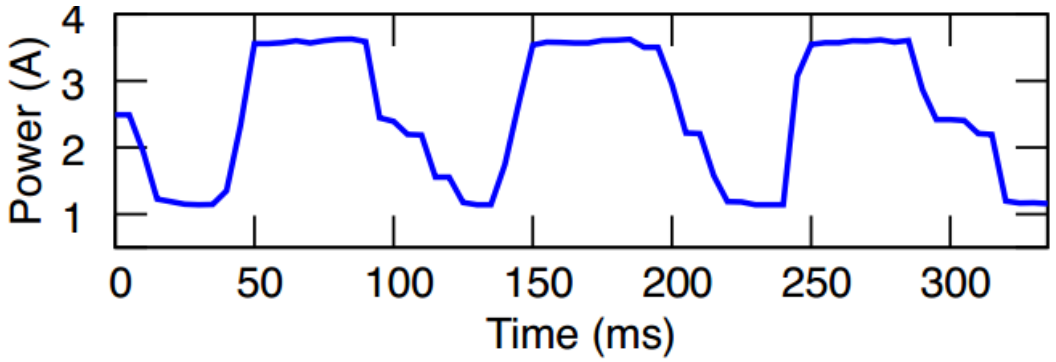

Covert-Channel Attack

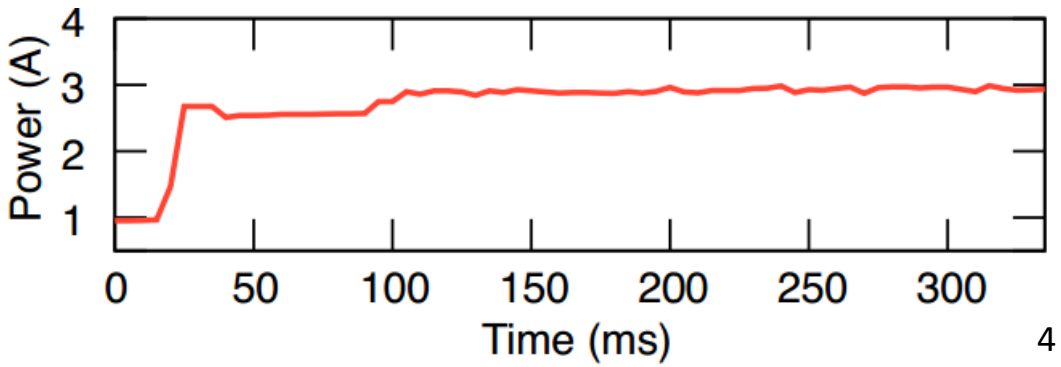

Challenge: Detection Operating Range?

- Complex systems
- Real workloads
- Range of attacks

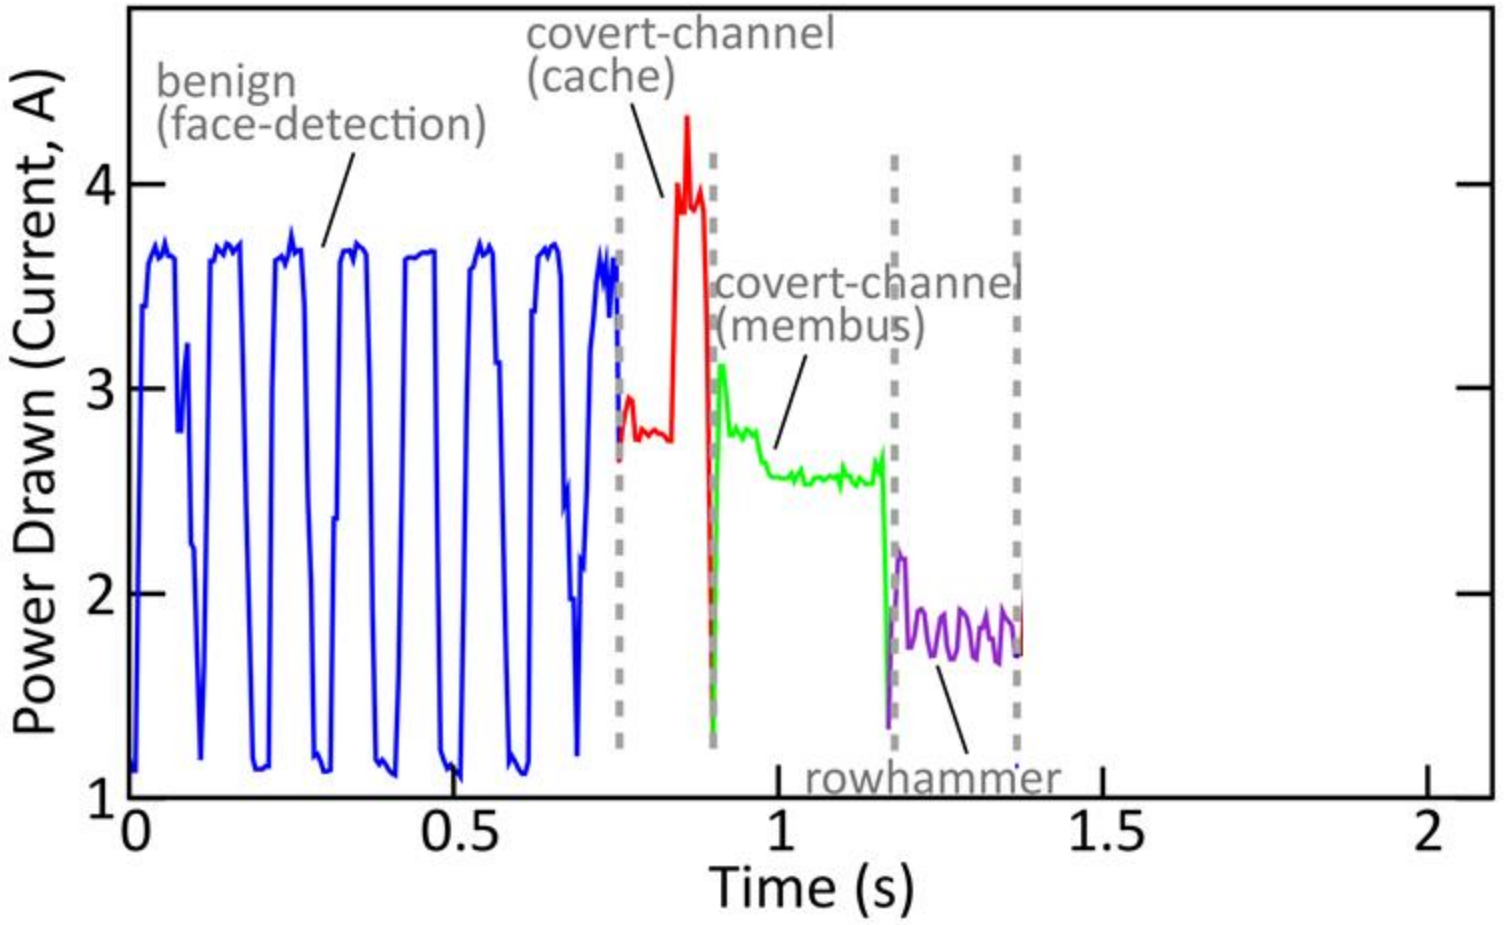

Challenge: Detection Operating Range?

- Complex systems
- Real workloads
- Range of attacks

- Evasive malware
  - Power-Mimicking

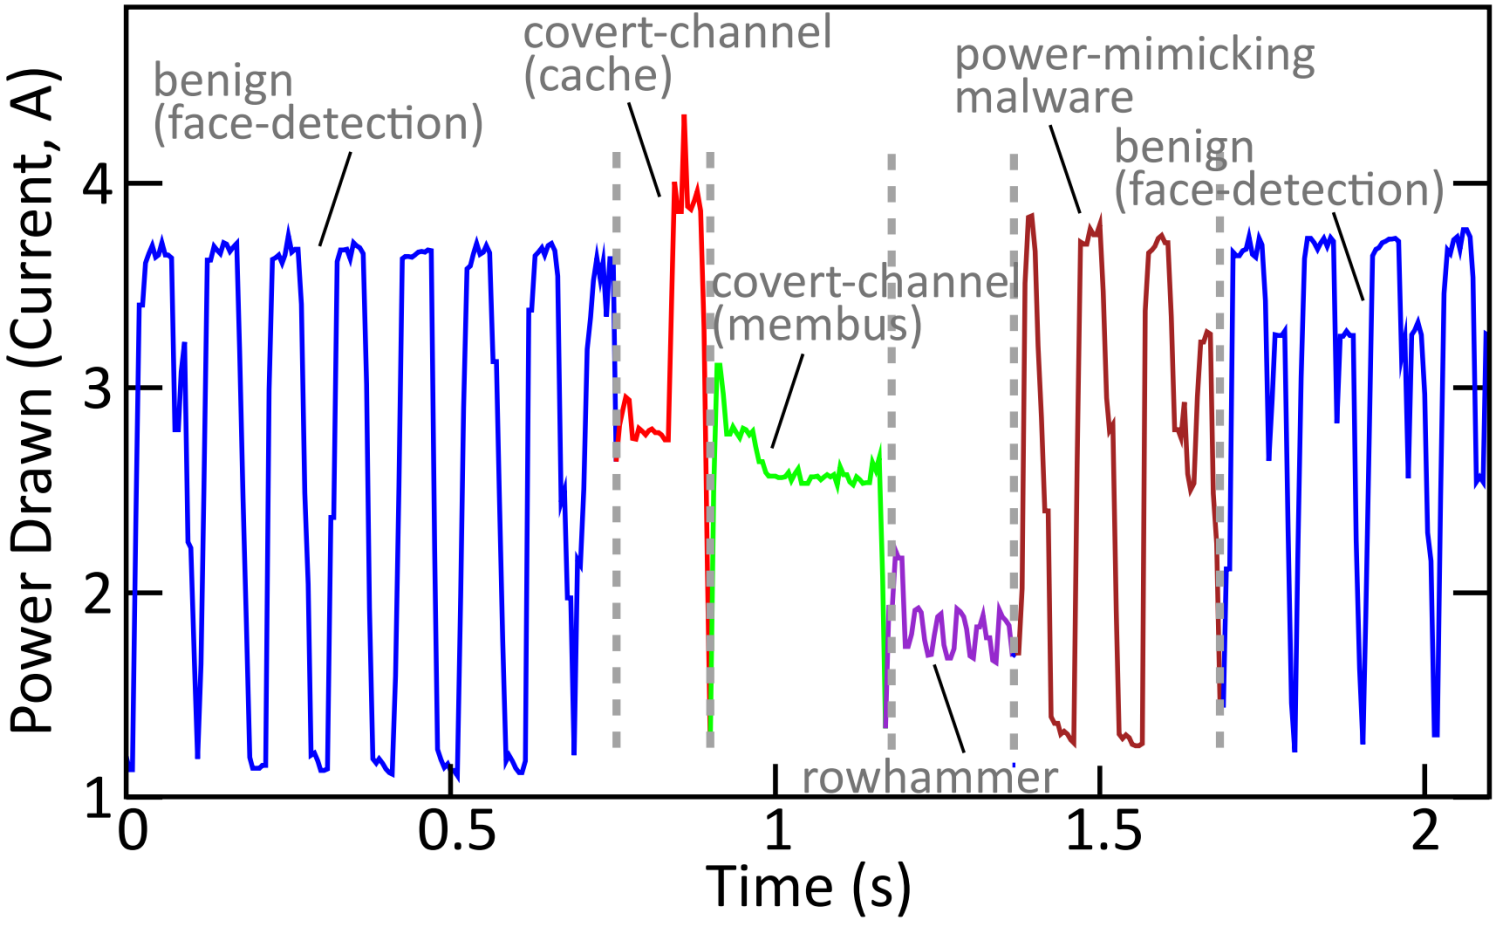

Challenge: Detection Operating Range?

- Complex systems
- Real workloads
- Range of attacks

- Evasive malware

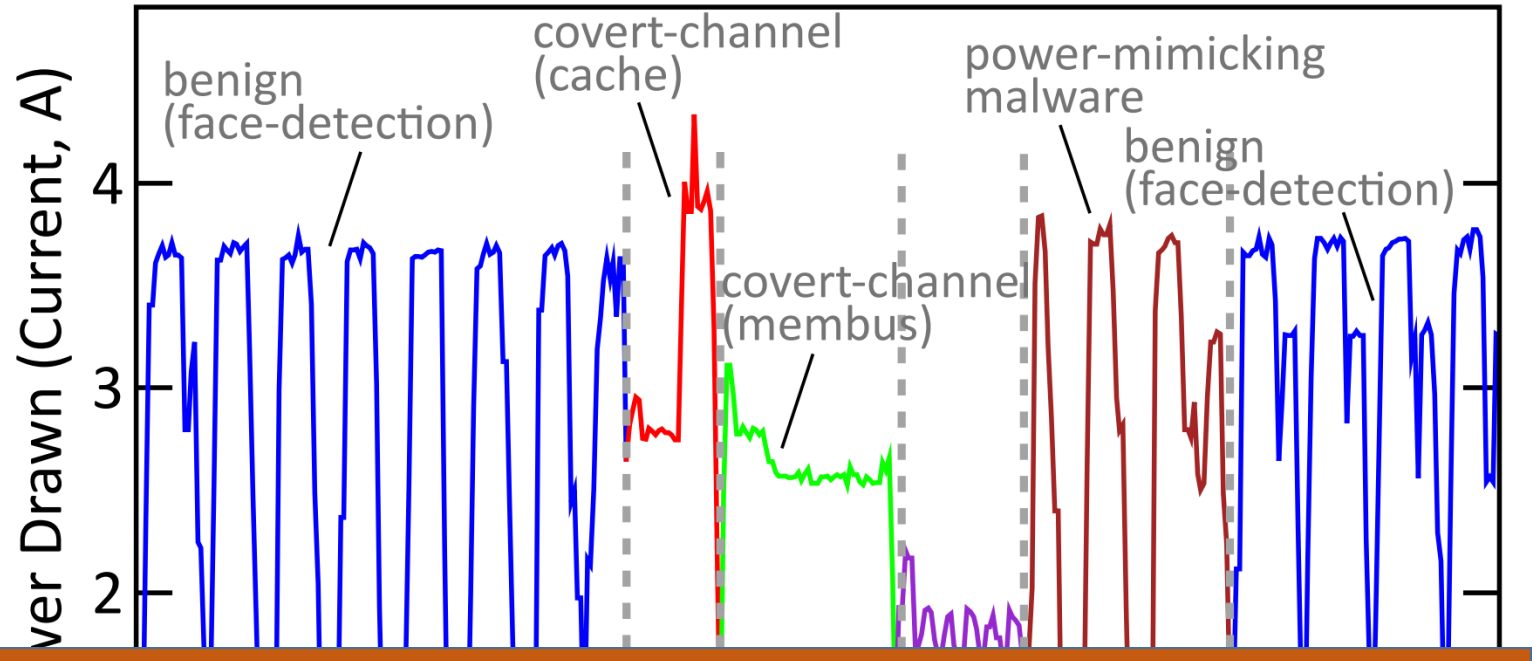

Rowhammer cannot evade
Covert channel and Spectre bandwidth reduced by 36x and 7x

Related Work in Power-Based Detectors

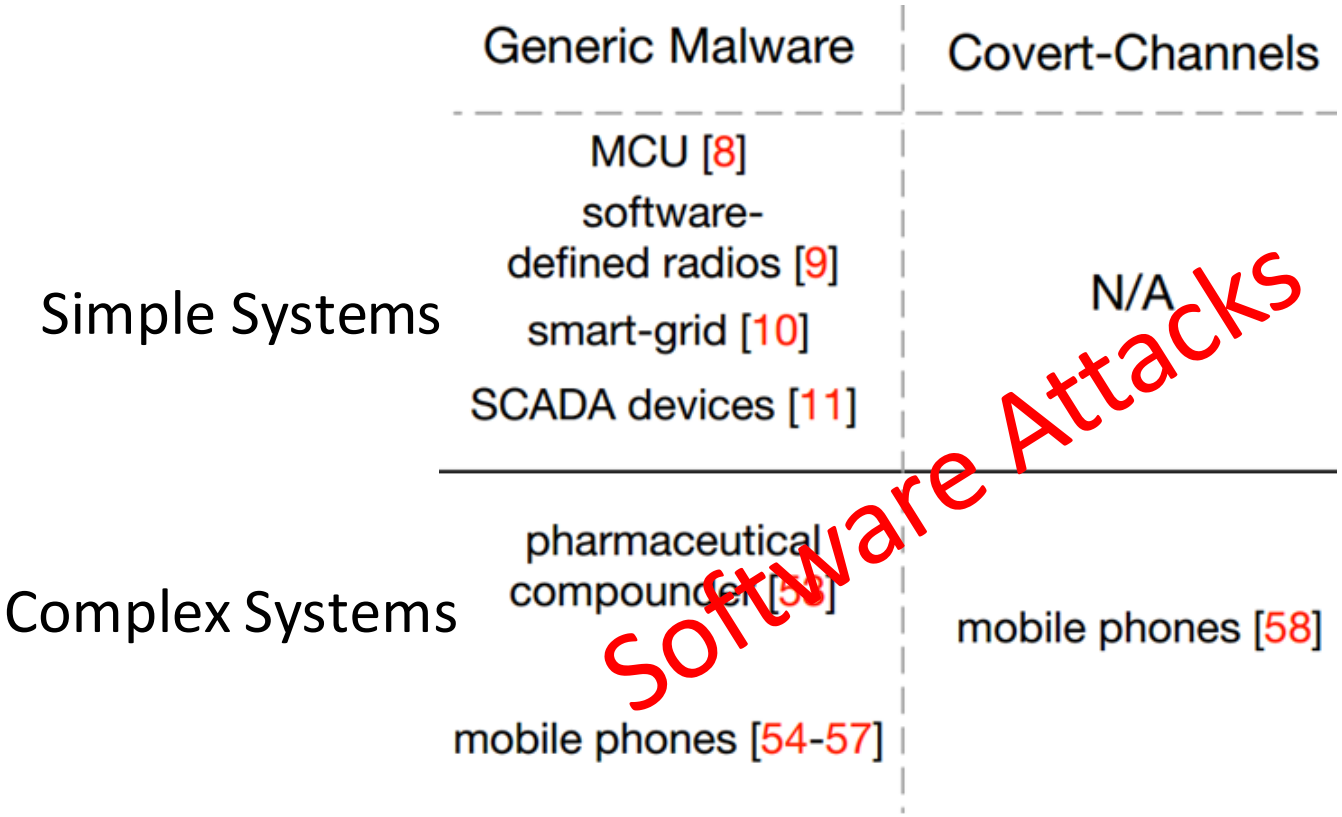

Related Work in Power-Based Detectors

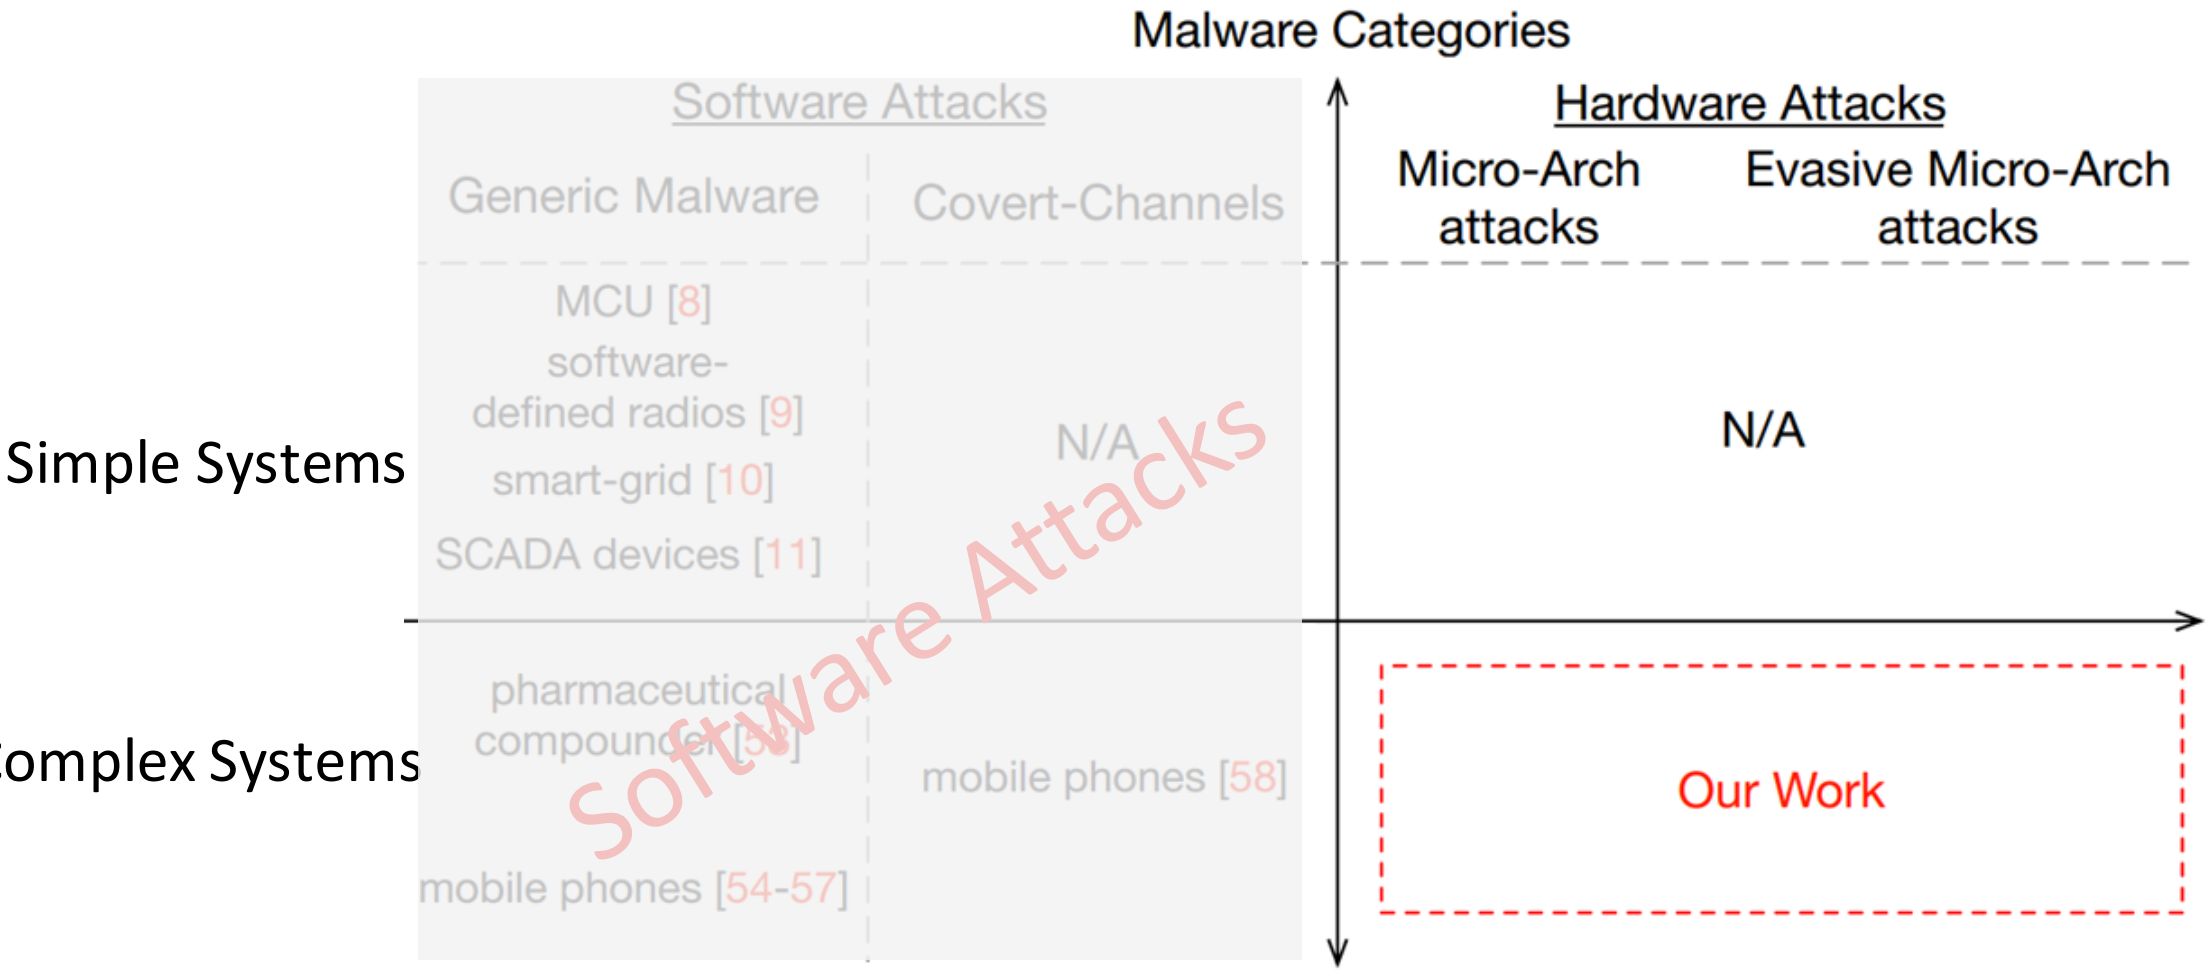

Synthesizing Evasive Power-Mimicking Attacks

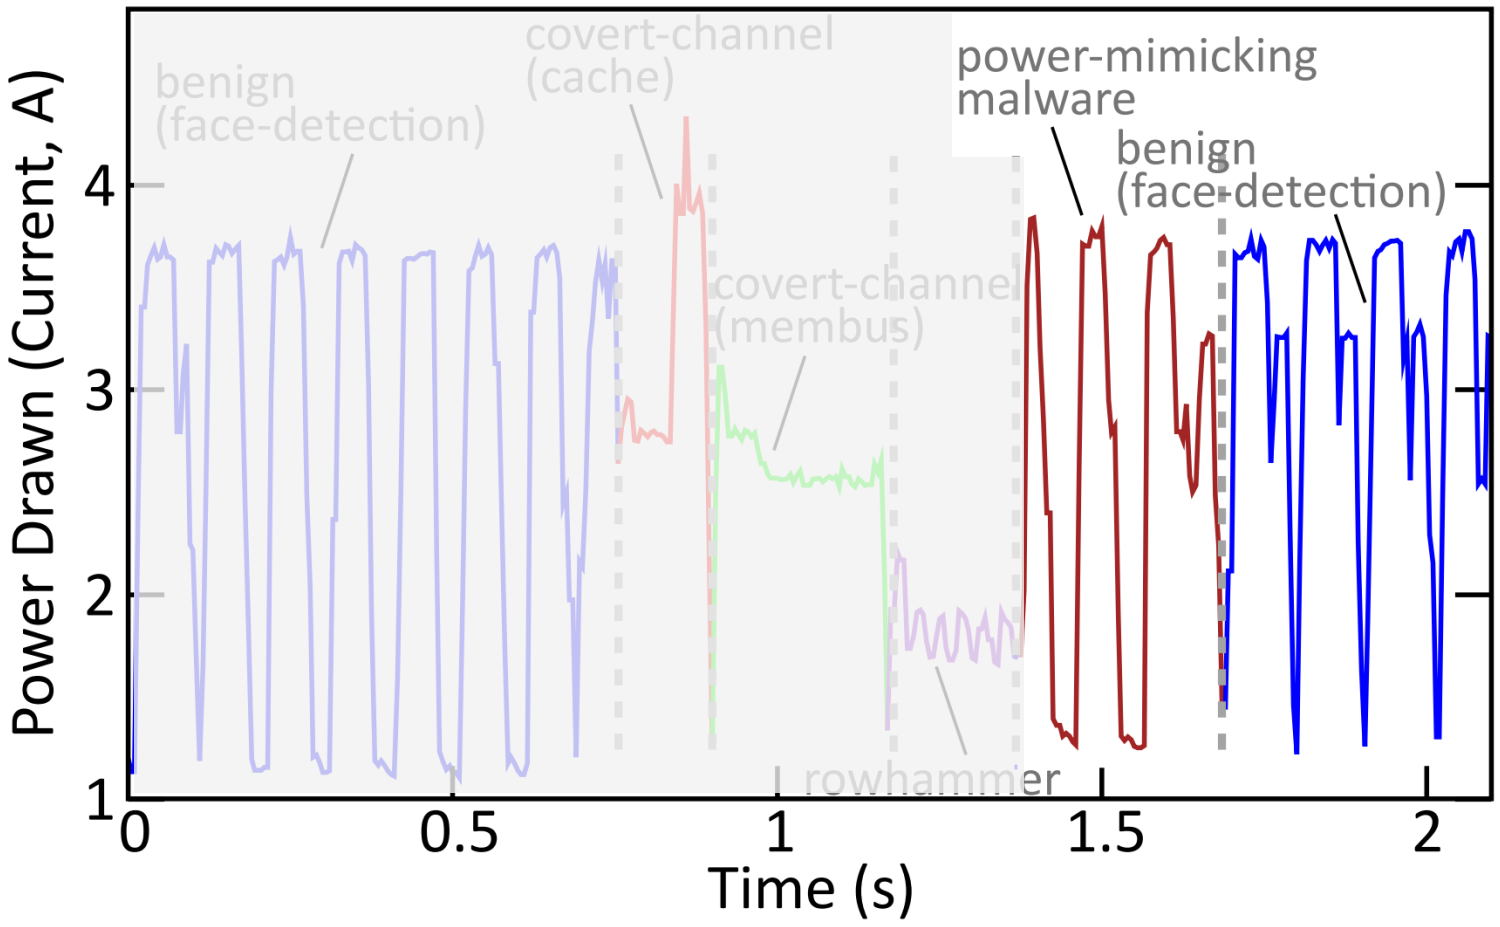

Synthesizing Evasive Power-Mimicking Attacks

Face Detection

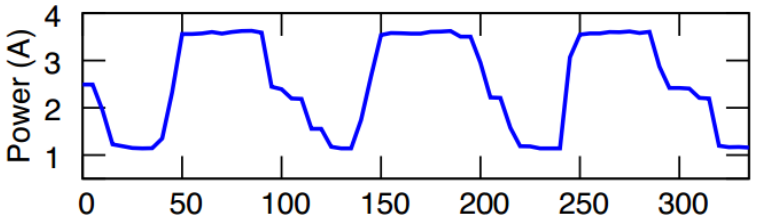

Covert-Channel, Before Mimicking

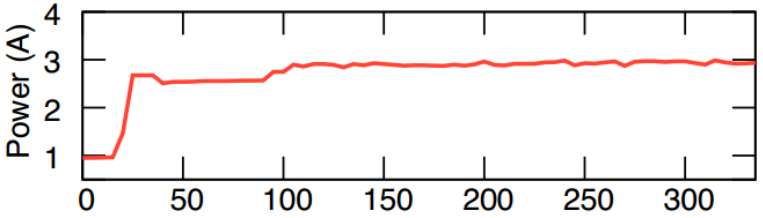

Power-Shaping Algorithm

Covert-Channel, After Mimicking

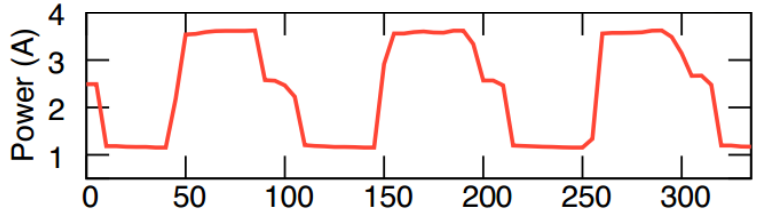

Synthesizing Evasive Power-Mimicking Attacks

- Threat model
  - Attacker is aware of detector
  - Can purchase instances

Synthesizing Evasive Power-Mimicking Attacks

- Threat model
  - Attacker is aware of detector
  - Can purchase instances
- Construct evasive attacks

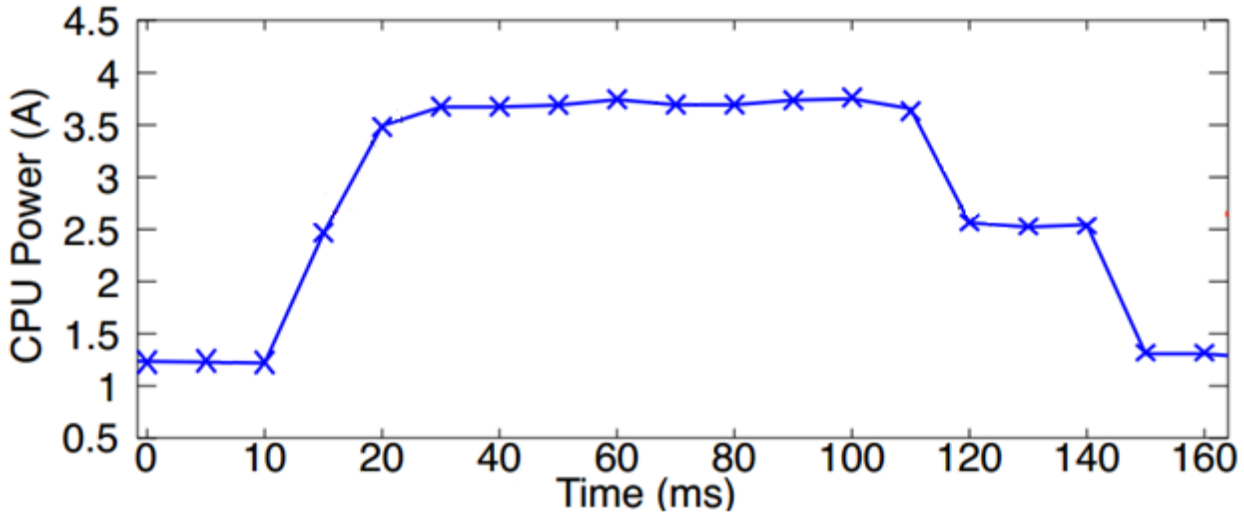

Synthesizing Evasive Power-Mimicking Attacks

- Threat model
  - Attacker is aware of detector
  - Can purchase instances
- Construct evasive attacks
  - Atomic tasks

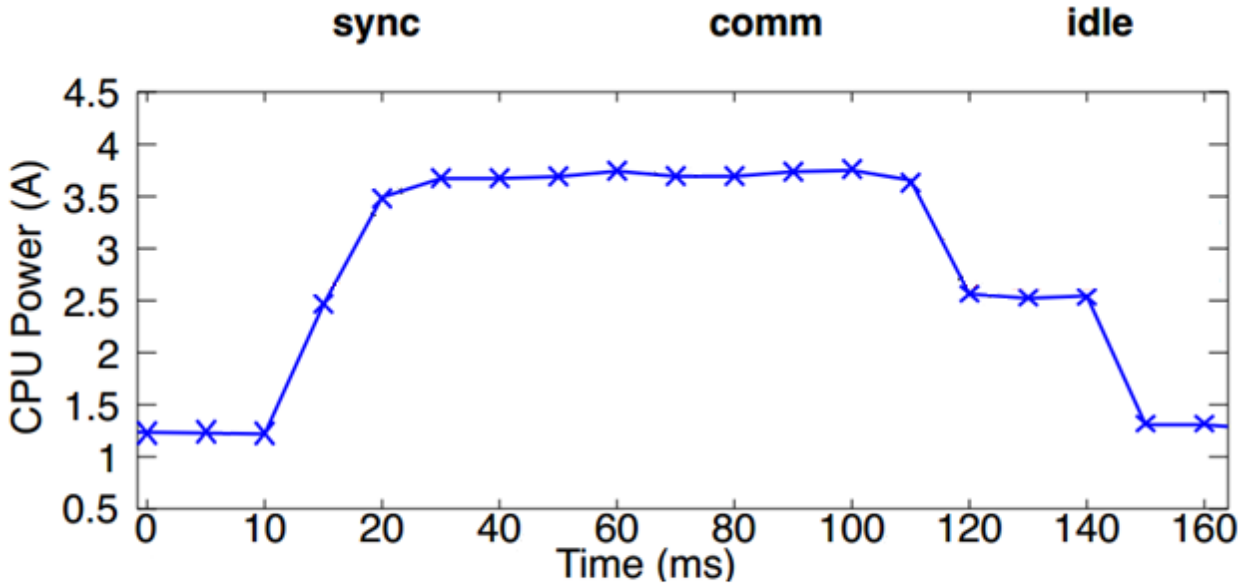

Synthesizing Evasive Power-Mimicking Attacks

- Threat model
  - Attacker is aware of detector
  - Can purchase instances
- Construct evasive attacks
  - Atomic tasks
  - Atomic task power

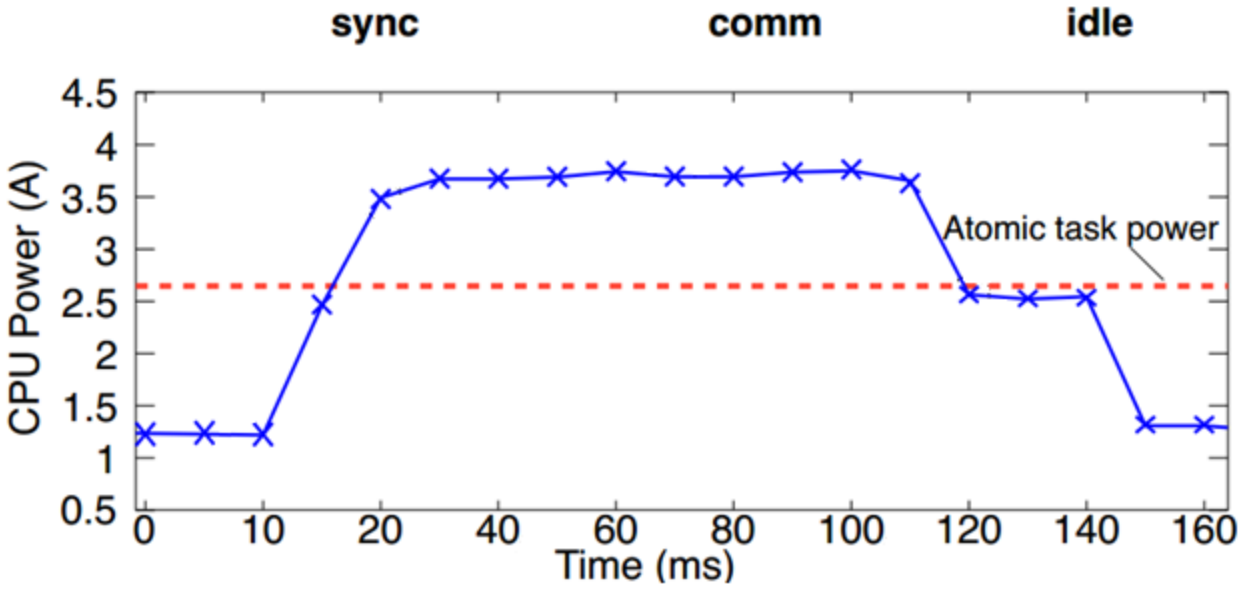

Synthesizing Evasive Power-Mimicking Attacks

- Threat model
  - Attacker is aware of detector
  - Can purchase instances
- Construct evasive attacks
  - Atomic tasks
  - Atomic task power
  - Inject tasks in basic time units
  - Shape power
    - SIMD, Cache etc.

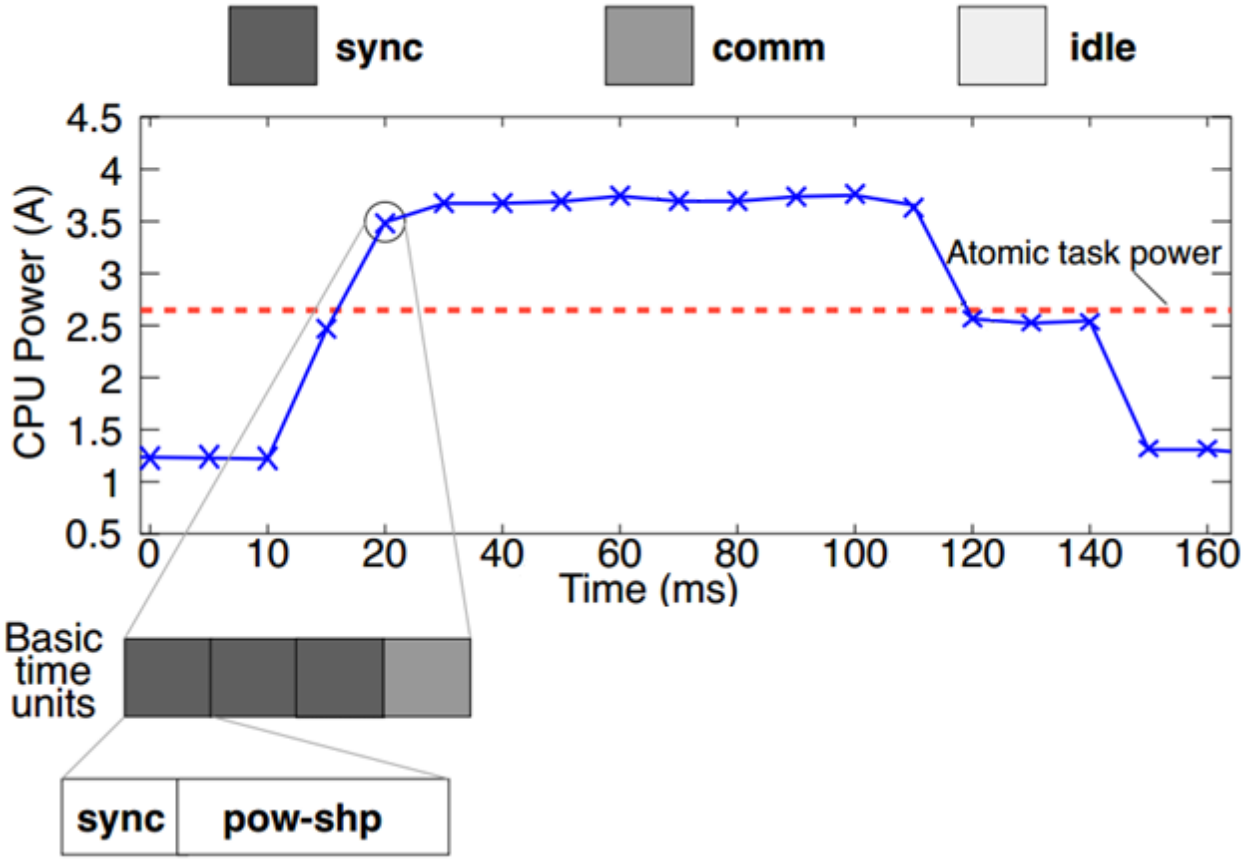

Synthesizing Evasive Power-Mimicking Attacks

- Threat model
  - Attacker is aware of detector
  - Can purchase instances
- Construct evasive attacks
  - Atomic tasks
  - Atomic task power
  - Inject tasks in basic time units
  - Shape power
    - SIMD, Cache etc.

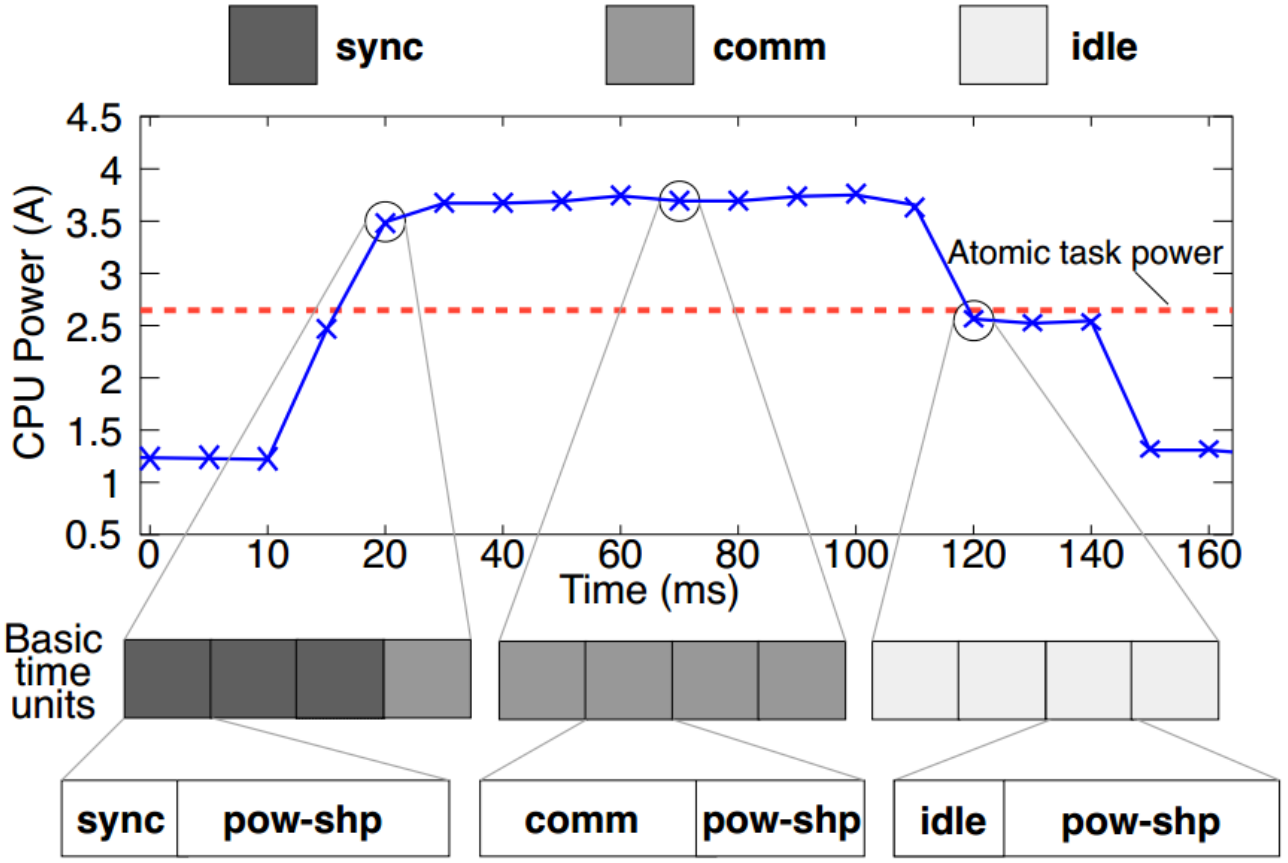

Evasive Power-Mimicking Attack Design Trade-Off

- Spend time to
  - Shape power?
  - OR perform a malicious task?

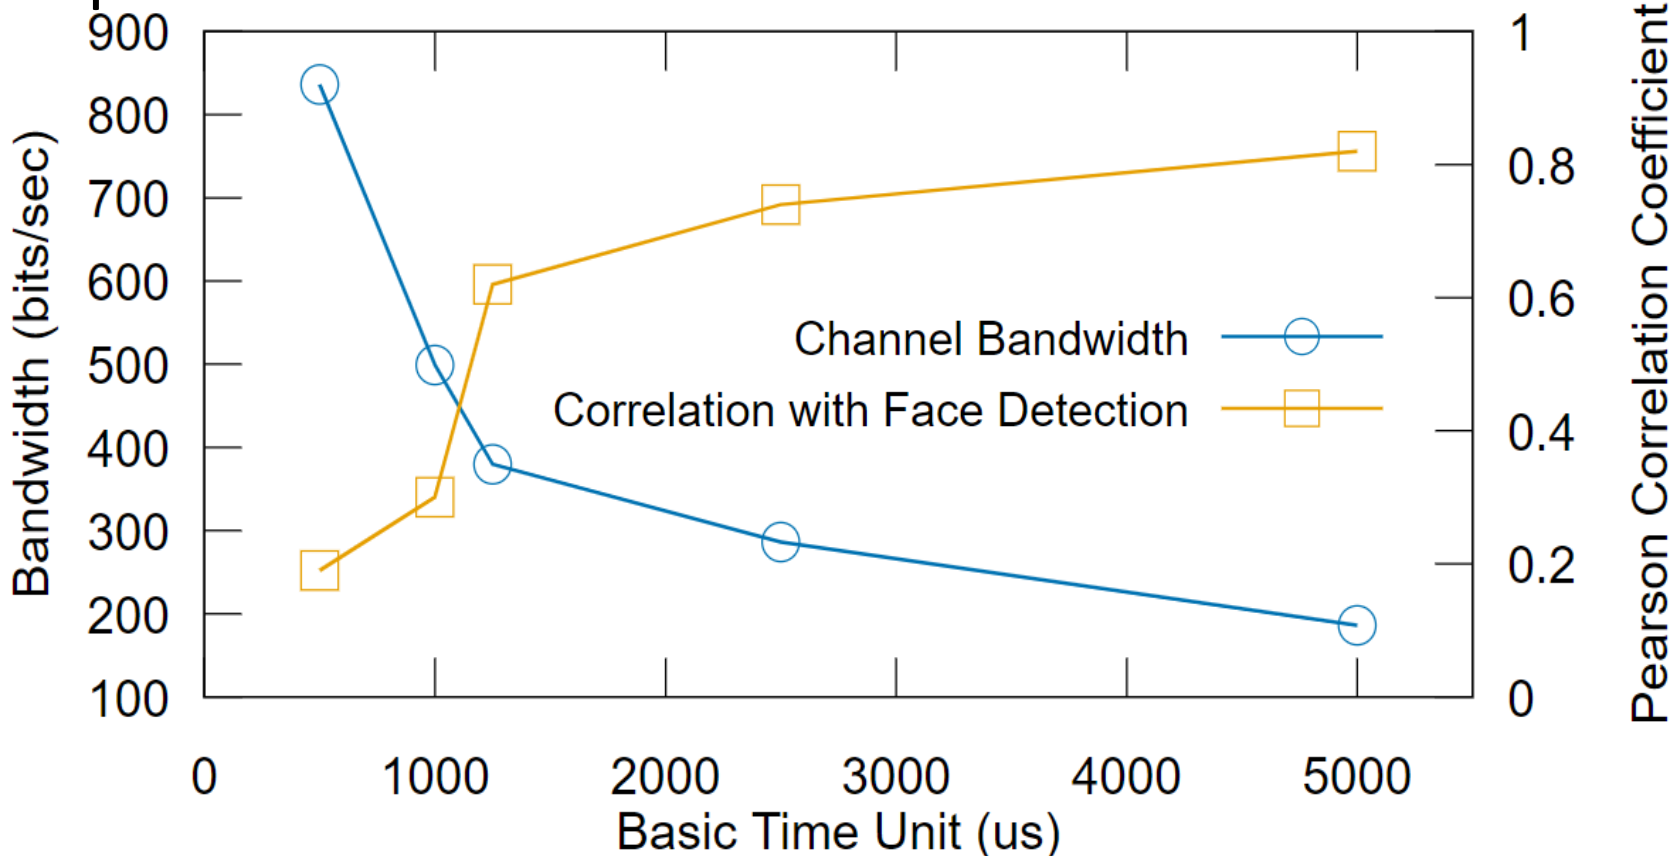

Power-Anomaly Detector: ML Training

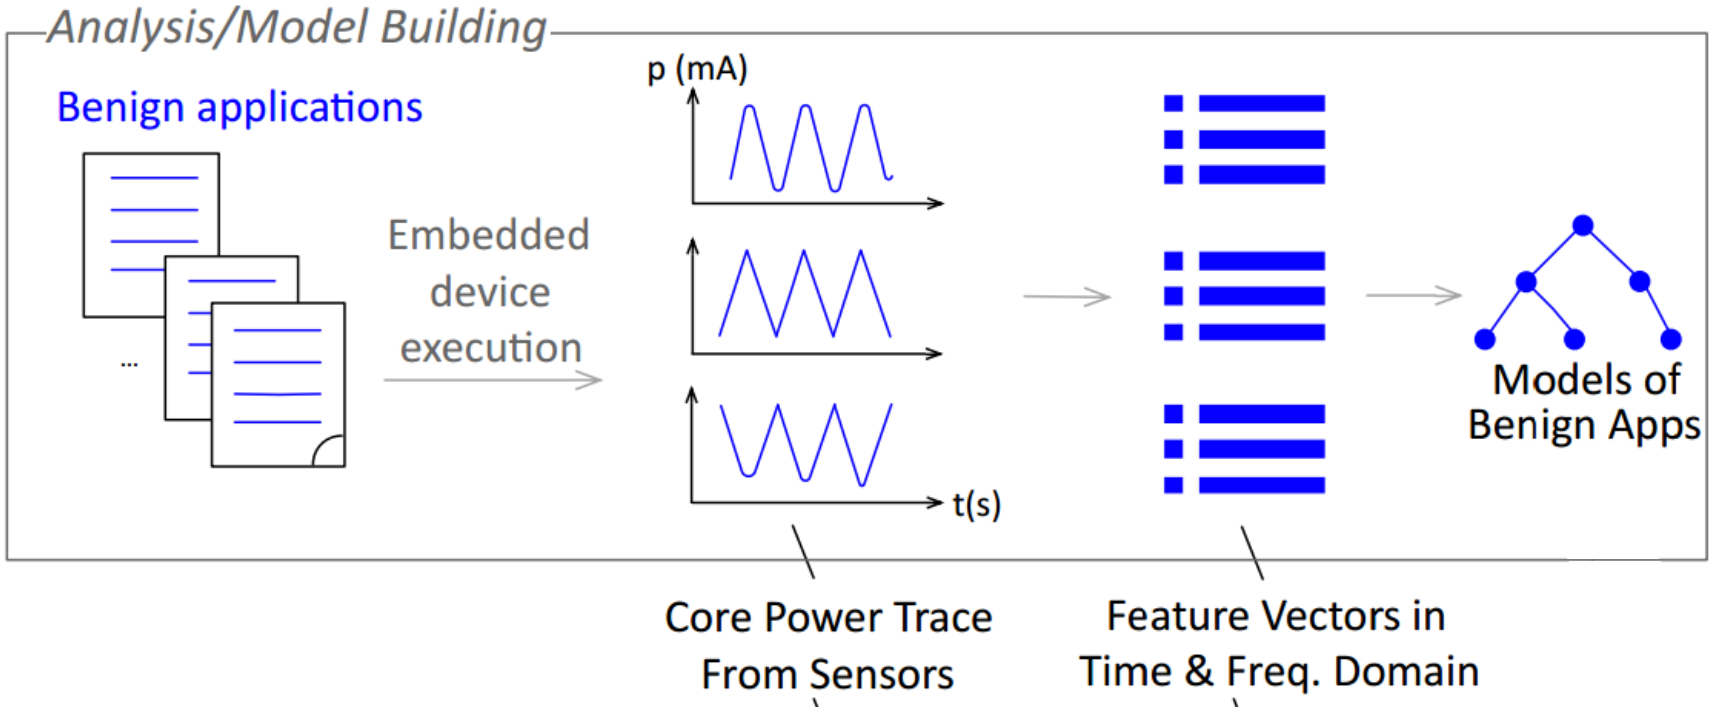

Power-Anomaly Detector: ML Testing

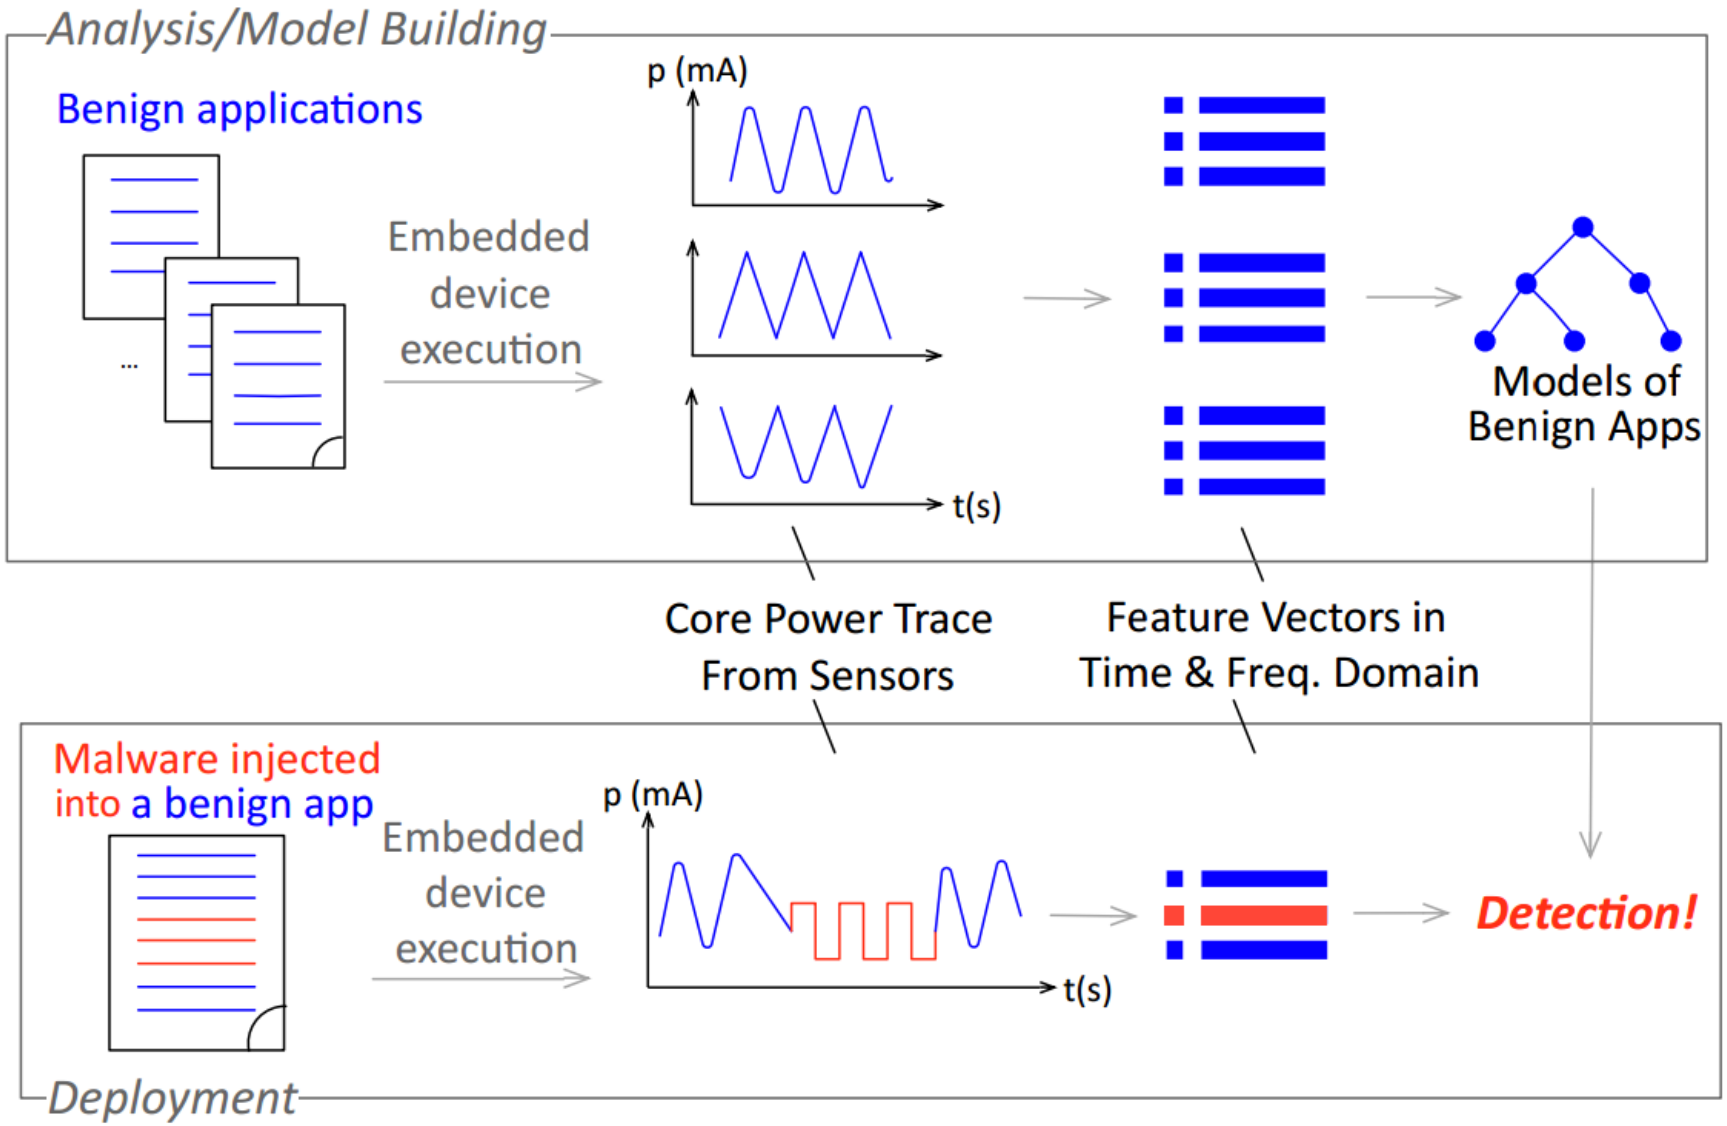

Evaluation Setup

- Benign applications
  - Computer vision, Crypto, Compression, Robotics
- Attacks
  - Cache/Memory covert-channel, Rowhammer, Spectre
- Lightweight Linux
  - Power constrained and compute optimized
- Odroid XU3
  - Onboard power sensor @up to 200HZ
- Machine learning models
  - SVM, LSTM

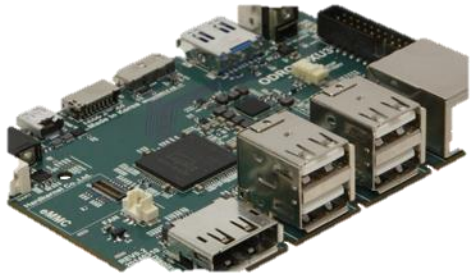

Evaluation: Benign Applications

Computer vision

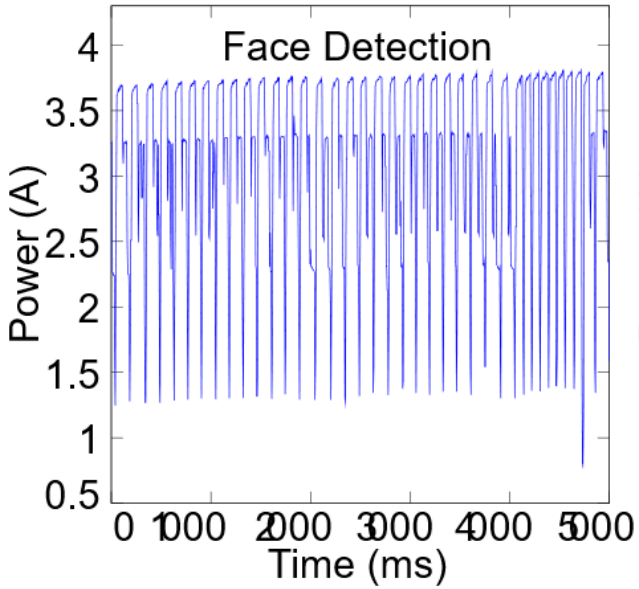

Crypto

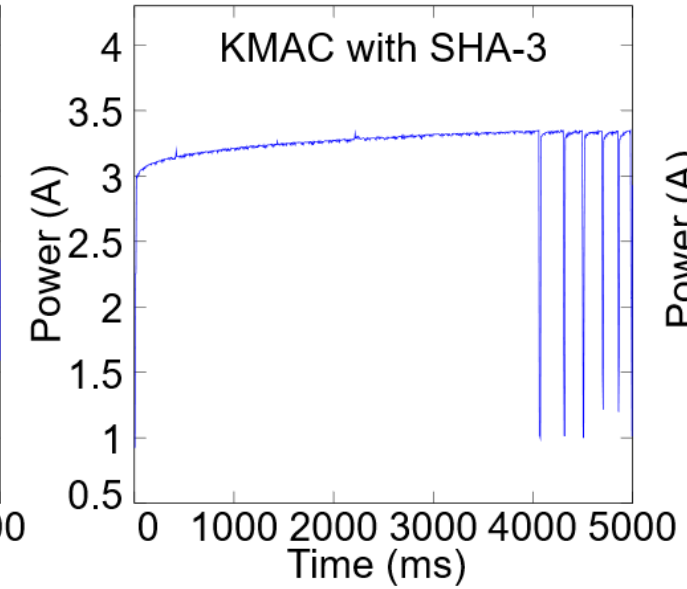

Compression

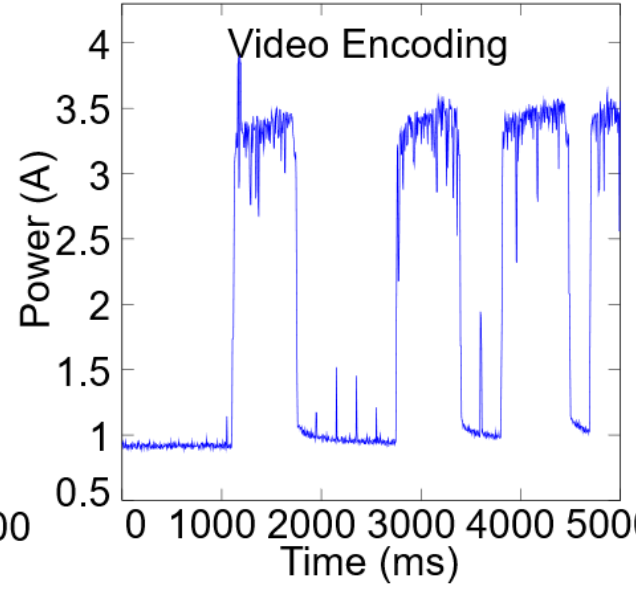

Robotics

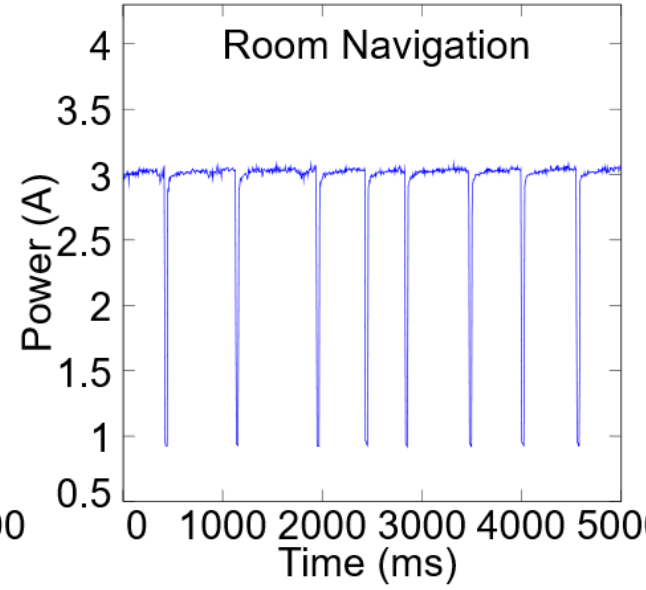

The University of Texas at Austin
Electrical and Computer Engineering
Cockrell School of Engineering

Using Power-Anomalies to Counter Evasive Micro-Architectural Attacks in Embedded Systems

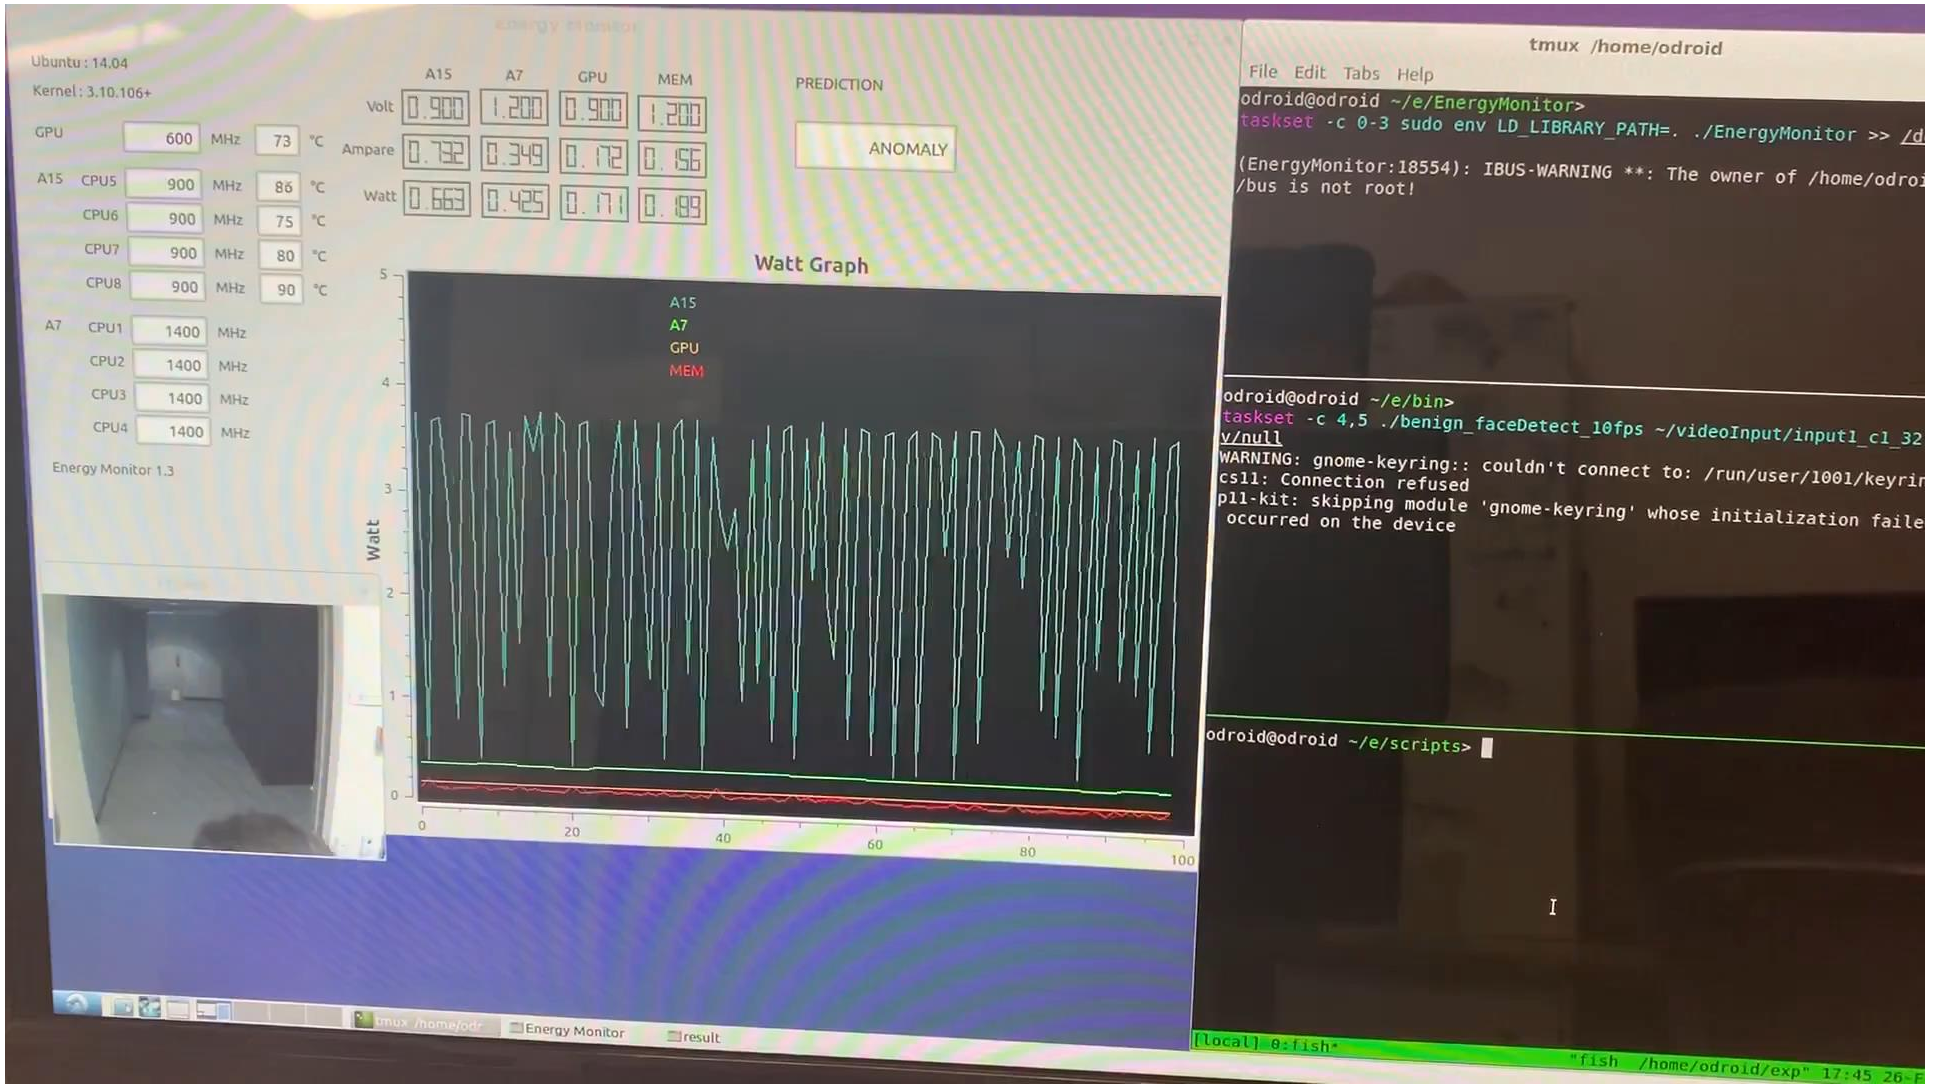

Results: Baseline Attacks

- Area Under Curve (AUC)

	One-Class SVM				
Benign / Attacks	Face Detection	Video Encoder	SHA-3 KMAC	Room Navigation	GeoMean
Cache Covert Channel	0.9964	0.9996	0.9880	0.9558	**0.9848**
Memory-bus Covert Channel	0.9938	1.0000	0.9814	0.9179	**0.9727**
Spectre	0.9980	0.9879	0.9763	0.9919	**0.9885**
Rowhammer	0.9928	0.9999	0.9777	0.9850	**0.9863**
GeoMean	**0.9927**	**0.9968**	**0.9828**	**0.9622**	

Results: Baseline Attacks

- Area Under Curve (AUC)

	One-Class SVM				
Benign / Attacks	Face Detection	Video Encoder	SHA-3 KMAC	Room Navigation	GeoMean
Cache Covert Channel	0.9964	0.9996	0.9880	0.9558	**0.9848**
Memory-bus Covert Channel	0.9938	1.0000	0.9814	0.9179	**0.9727**
Spectre	0.9980	0.9879	0.9763	0.9919	**0.9885**
Rowhammer	0.9928	0.9999	0.9777	0.9850	**0.9863**
GeoMean	**0.9927**	**0.9968**	**0.9828**	**0.9622**	

	LSTM					
Benign / Attacks	Face Detection	Video Encoder	SHA-3 KMAC	Room Navigation	All Benign Apps	GeoMean
Cache Covert Channel	0.9985	0.9847	1.0000	0.9615	**0.9917**	**0.9861**
Memory-bus Covert Channel	0.9969	0.9816	0.9999	0.9589	**0.9898**	**0.9842**
Spectre	0.9974	0.9821	0.9939	0.9956	**0.9962**	**0.9822**
Rowhammer	0.9972	0.9828	1.0000	0.9823	**0.9898**	**0.9906**
GeoMean	**0.9975**	**0.9828**	**0.9984**	**0.9745**	**0.9919**	

Results: Evasive Attacks

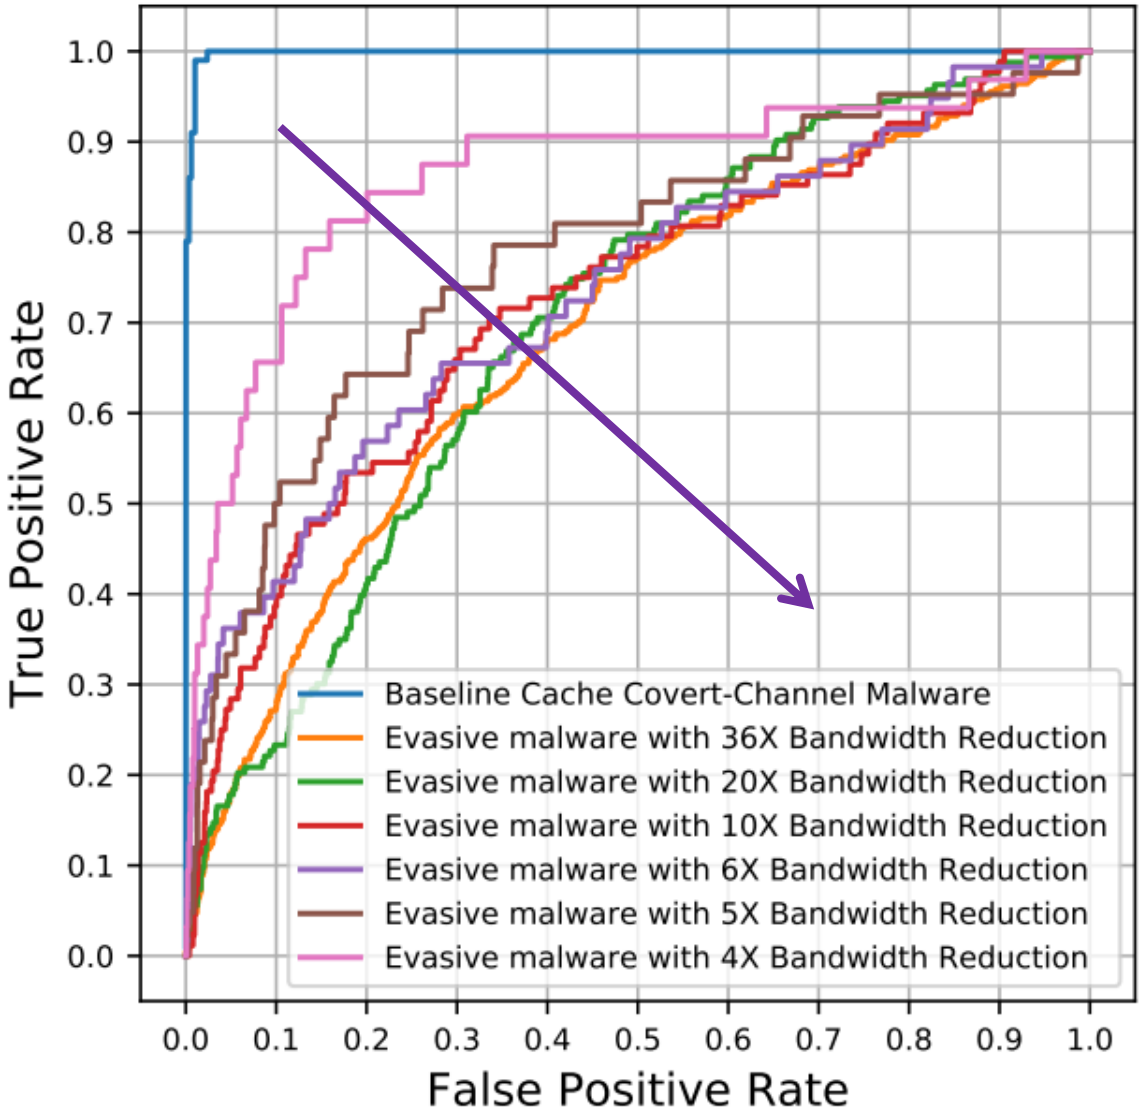

- Receiver Operating Characteristic (ROC) curve
- Mimic face detection

Results: Evasive Attacks

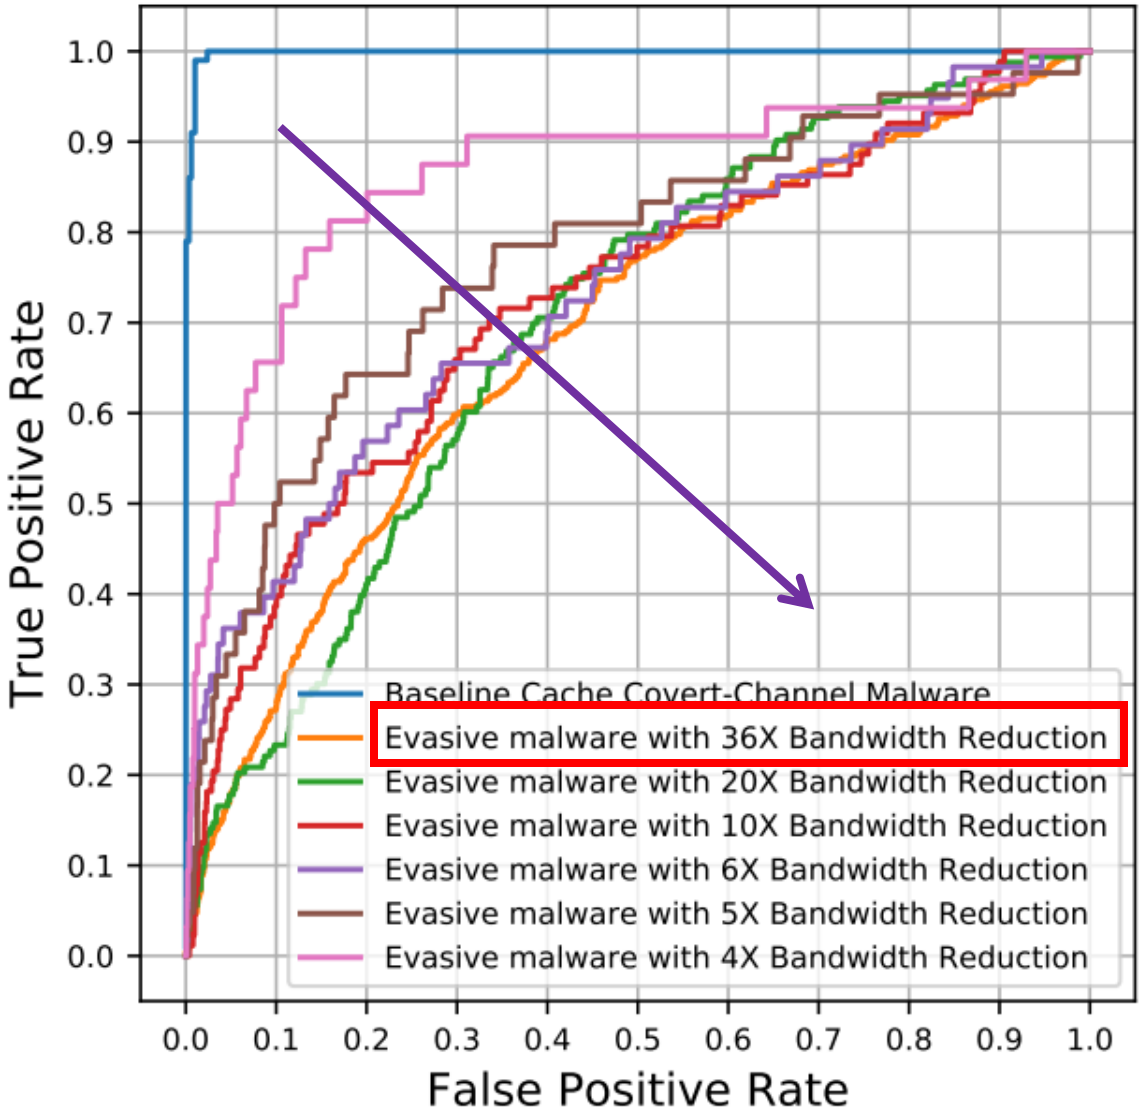

- Receiver Operating Characteristic (ROC) curve
- Mimic face detection
- Rowhammer cannot evade
  - Can shape power
  - Failed to flip bits
- Spectre is still detected
  - Bandwidth drop from 7 to 1bps
  - Detected with higher false positive

Deployment Example: Smart Battery

- Existing simple compute resource
  - Nios II microcontroller on Altera Cyclone® IV
- Out-of-Band component
  - Plug-n-Play
- Smaller trust compute base
- Coarse-Grained measurements
  - Cheap

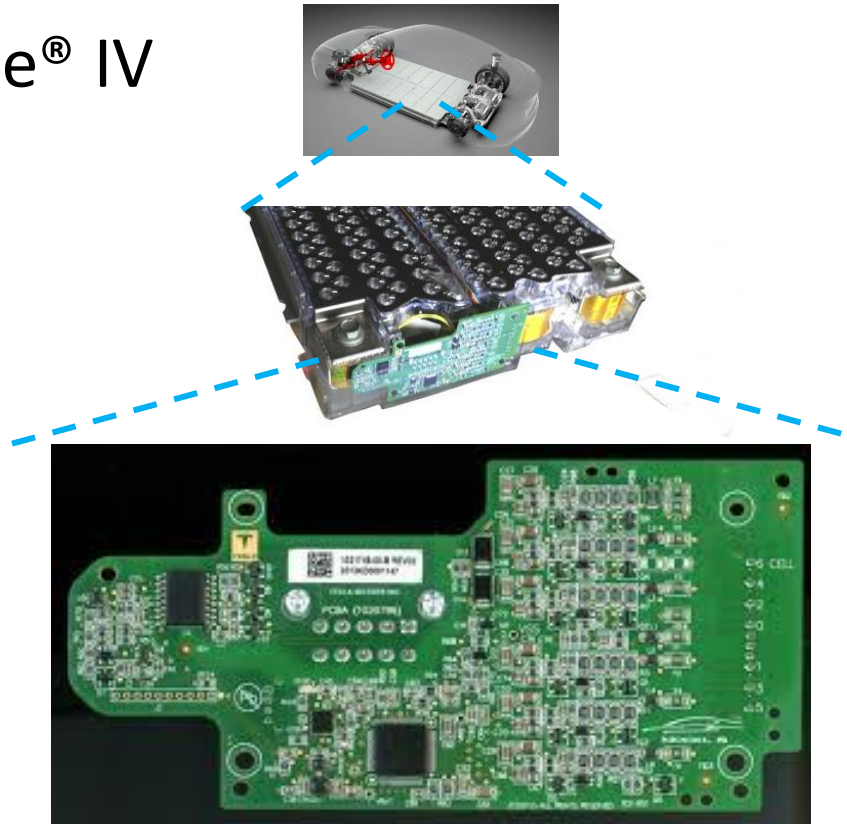

Summary

- Simple, out-of-band power-anomaly detector against recent microarchitectural attacks
- Introduces the first method to construct power-mimicking attacks
- Quantifies the operating range of machine learning power-anomaly detectors
- Demonstrates the tradeoff between malware stealth and damage

The University of Texas at Austin
Electrical and Computer Engineering
Cockrell School of Engineering

Using Power-Anomalies to Counter Evasive Micro-Architectural Attacks in Embedded Systems

Thank You!

http://spark.ece.utexas.edu/

Using Power-Anomalies to Counter Evasive Micro-Architectural Attacks in Embedded Systems

BACKUP SLIDES