NIGHTs-WATCH

A Cache-Based Side-Channel Intrusion Detector using Hardware Performance Counters

Maria Mushtaq, Ayaz Akram, Khurram Bhatti, Maham Chaudhry, Vianney Lapotre, Guy Gogniat

Contact: khurram.bhatti@itu.edu.pk

Outline

- Cache-based Side-Channel Attacks
  - The Memory Footprint
  - Hardware Performance Counters (HPCs)
  - Machine Learning Models
- NIGHTs-WATCH
  - The Big Picture
  - Selected HPCs & Machine Learning Models
- Experimental Results
  - Case Study-I: Flush+Reload on RSA
  - Case Study-II: Flush+Flush on AES

Motivation

- SCA defenses (mostly) offer *all-weather* protection & (often) heavily trade-off performance for protection
- Need-based protection could help –but accurate and fast need assessment is crucial
- Detection can be a first line of defense!

Side-Channel Attacks

- Cache-based SCAs exploit memory footprint of the victim process
  - Memory Access Timing: Can reveal from where the data/ instructions are being accessed
  - Memory Access Pattern: Can reveal what exactly is being processed

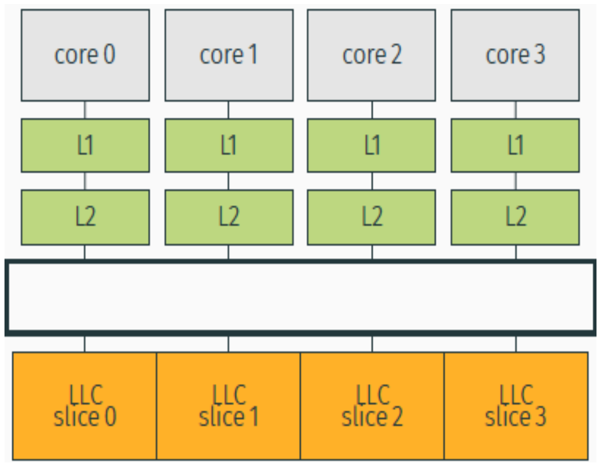

Figure: Courtesy Yarom et al.

Side-Channel Attacks

- Memory Access Timing Information

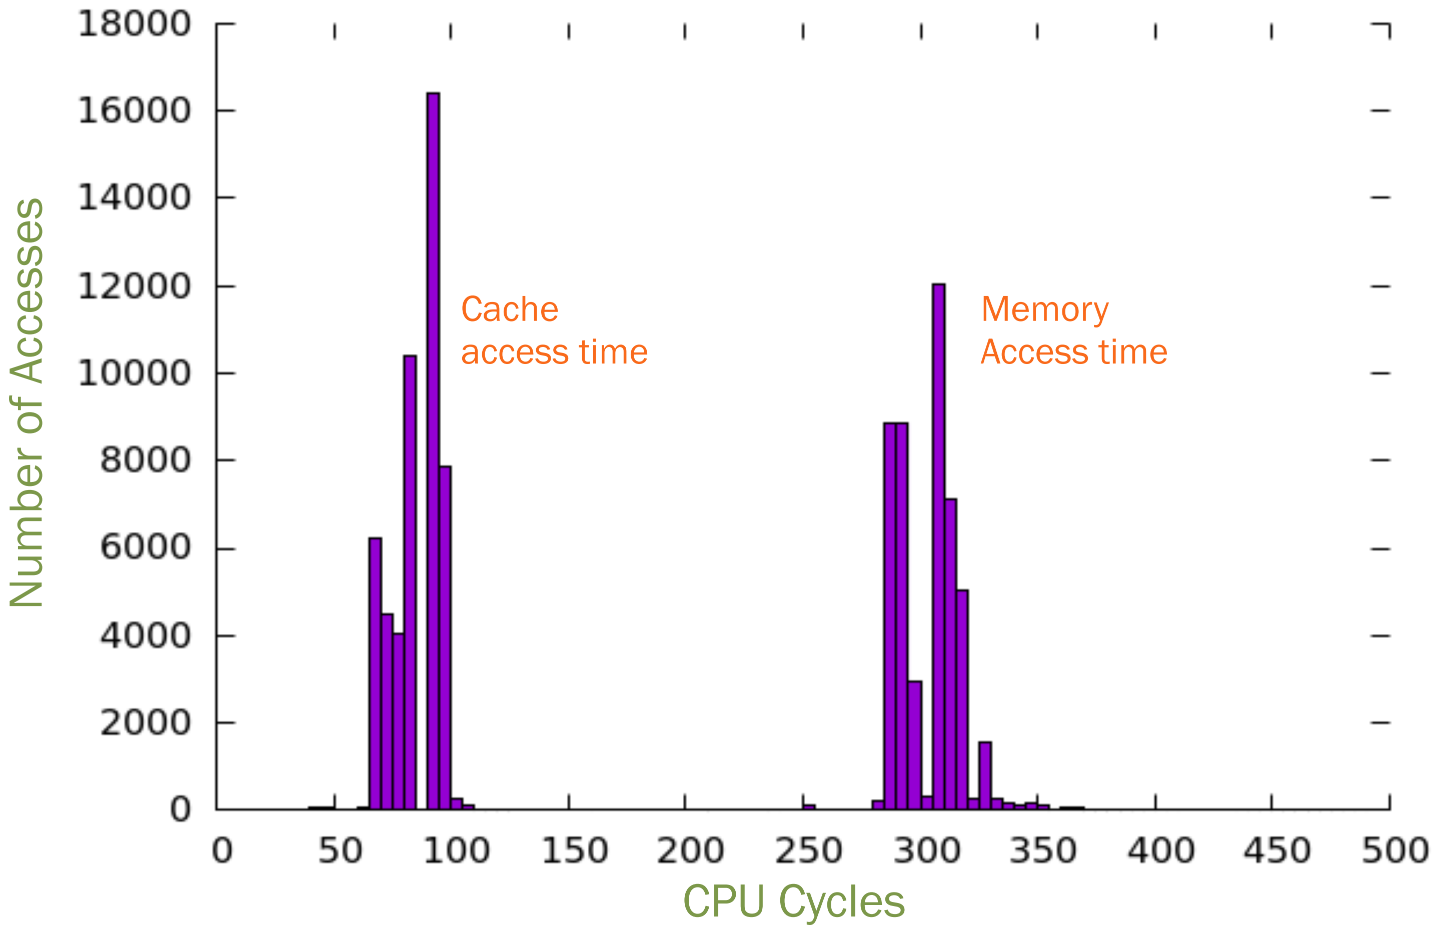

Results measured on Intel i5 for F+R Attack implementation.

Side-Channel Attacks

- Memory Access Pattern Information

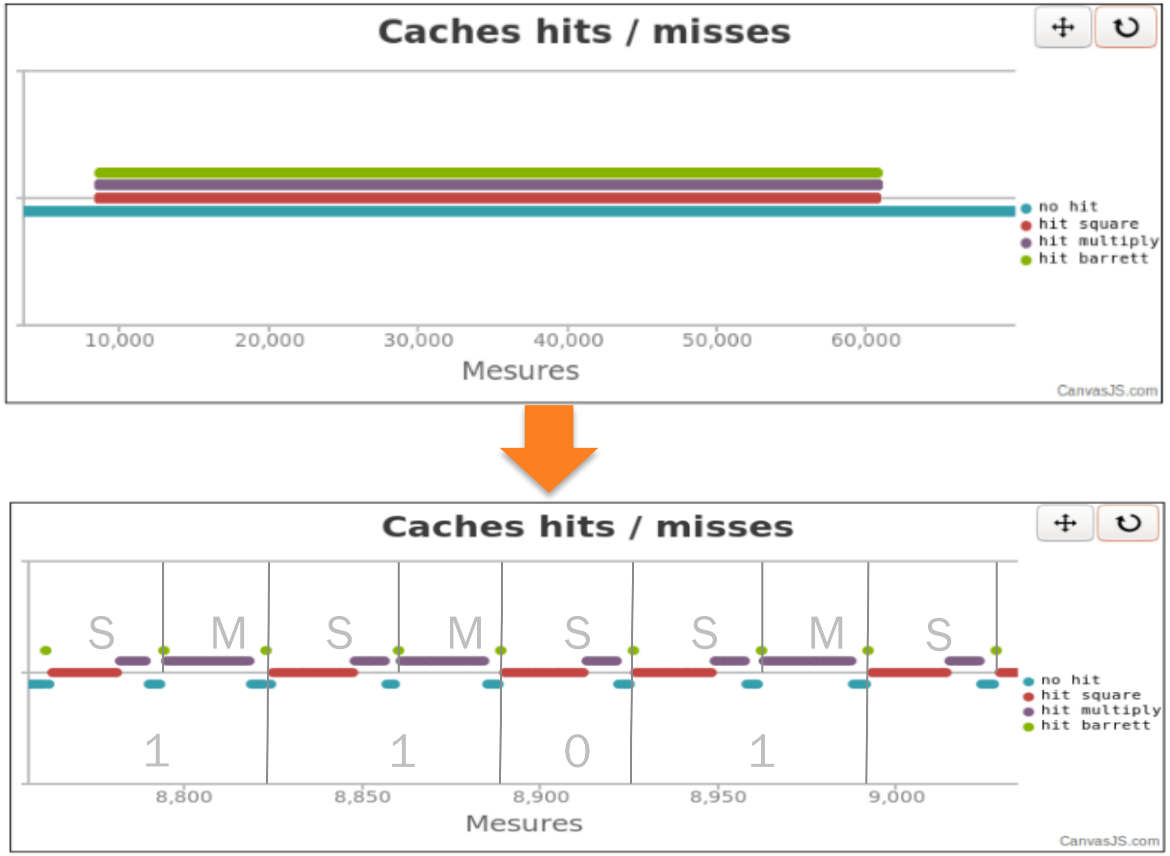

Results measured on i7 for F+R attack on RSA: Cache hit pattern for Square, Multiply, and Barrett operations.

SCA Detection

- Attacks have their own memory footprint too!

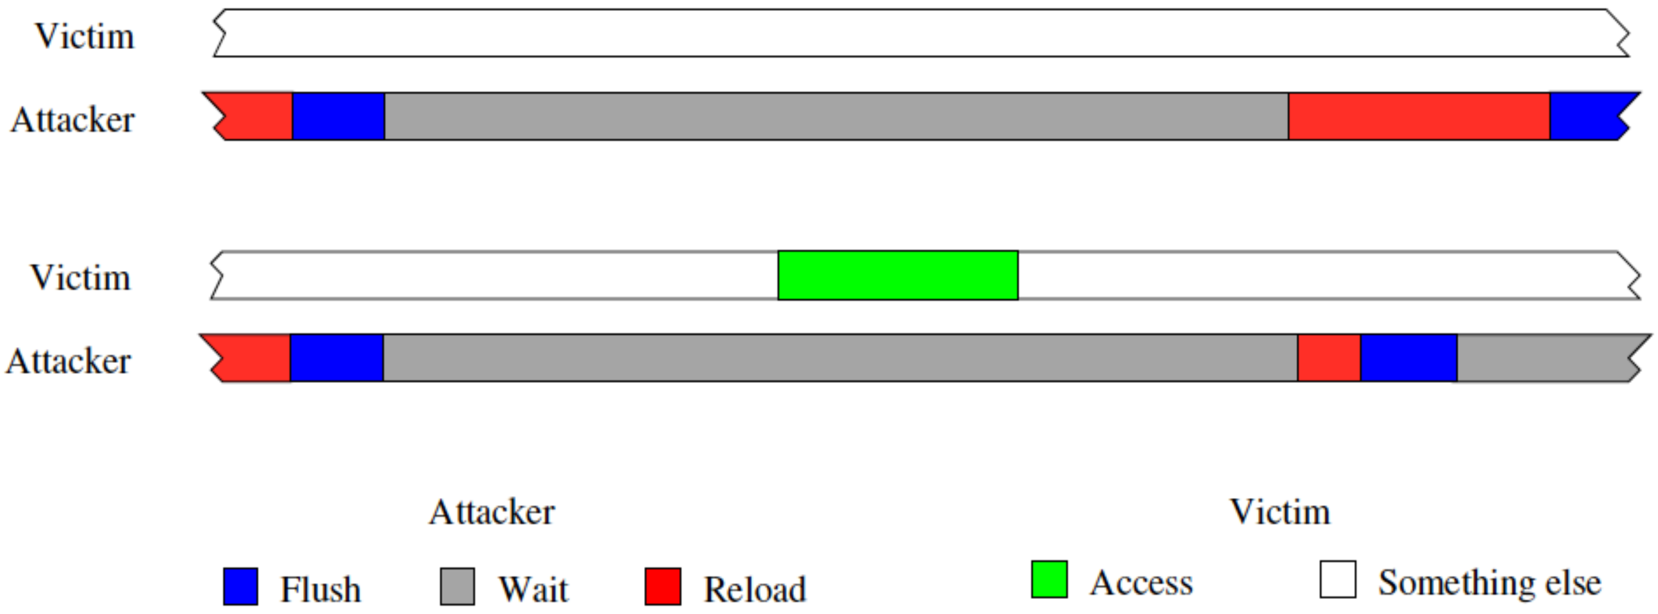

3-Phases of F+R Attack implementation

Figure: Courtesy Yarom et al.

SCA Detection

- Hardware Performance Counters (HPCs)
  - Specialized HW registers for performance monitoring that reveal run-time behavioral information of software

#	Scope	Hardware Performance Counters
1	Cache Level 1	Data Cache Misses (L1-DCM)
2		Instruction Cache misses (L1-ICM)
3		Total cache misses (L1-TCM)
4	Cache Level 2	Instruction cache accesses (L2-ICA)
5		Instruction Cache misses (L2-ICM)
6		Total Cache accesses (L2-TCA)
7		Total cache misses (L2-TCM)
8	Cache Level 3	Instruction cache accesses (L3-ICA)
9		Total Cache accesses (L3-TCA)
10		Total cache misses (L3-TCM)
11	System-wide	Branch Miss Prediction (BR_MSP)
12		Total CPU Cycles (TOT_CYC)

SCA Detection

- Hardware Performance Counters (HPCs)

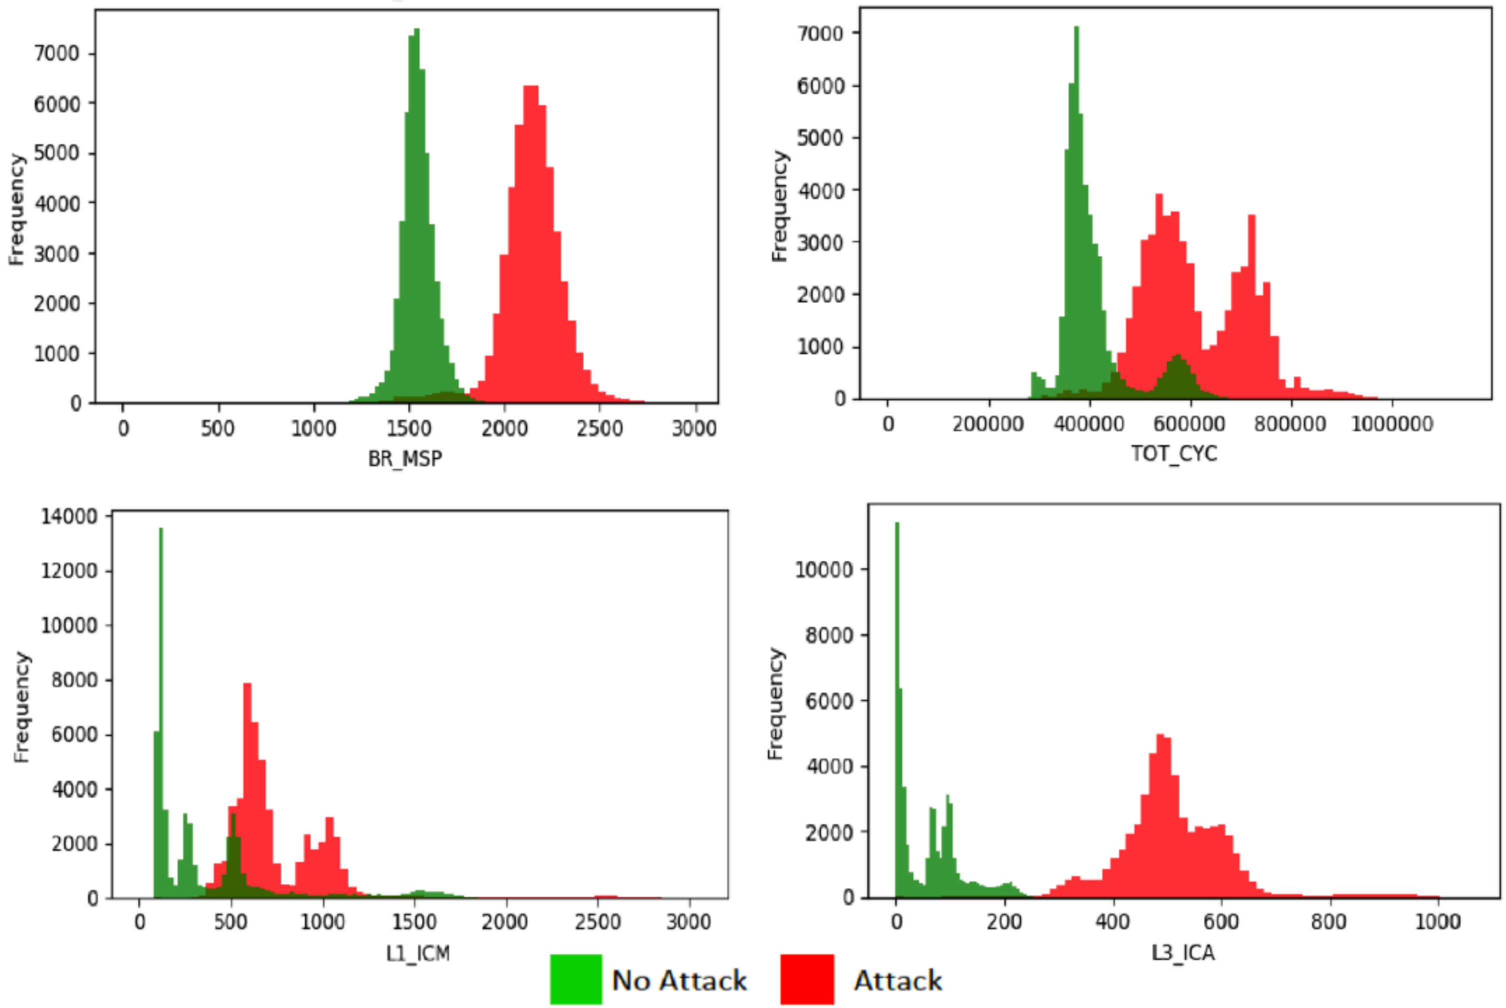

SCA Detection

- Hardware Performance Counters (HPCs)

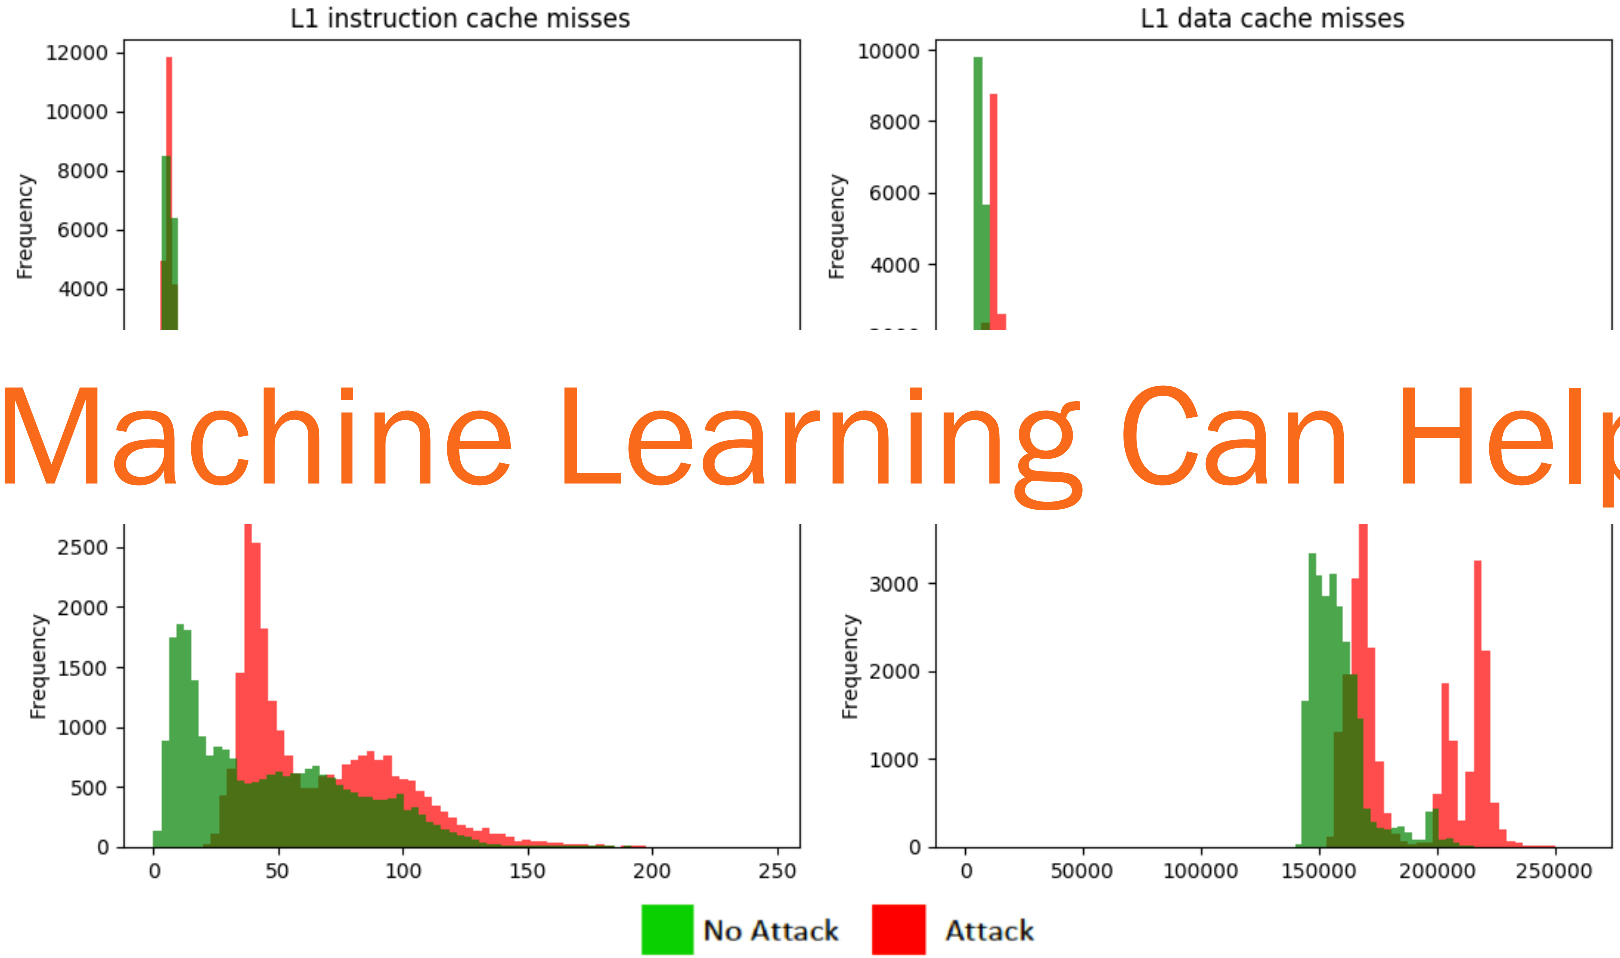

Machine Learning Can Help!

Outline

- Cache-based Side-Channel Attacks
  - The Memory Footprint
  - Hardware Performance Counters (HPCs)
  - Machine Learning Models
- NIGHTs-WATCH
  - The Big Picture
  - Selected HPCs & Machine Learning Models
- Experimental Results
  - Case Study-I: Flush+Reload on RSA
  - Case Study-II: Flush+Flush on AES

NIGHTs-WATCH

○ The Big Picture

Variable Load Conditions
Selected HPCs

Selected ML Models
One-Time Training Process
Scalable Set of ML Models

Trained ML Classifiers
Real-time HPC Data
Run-time Classification

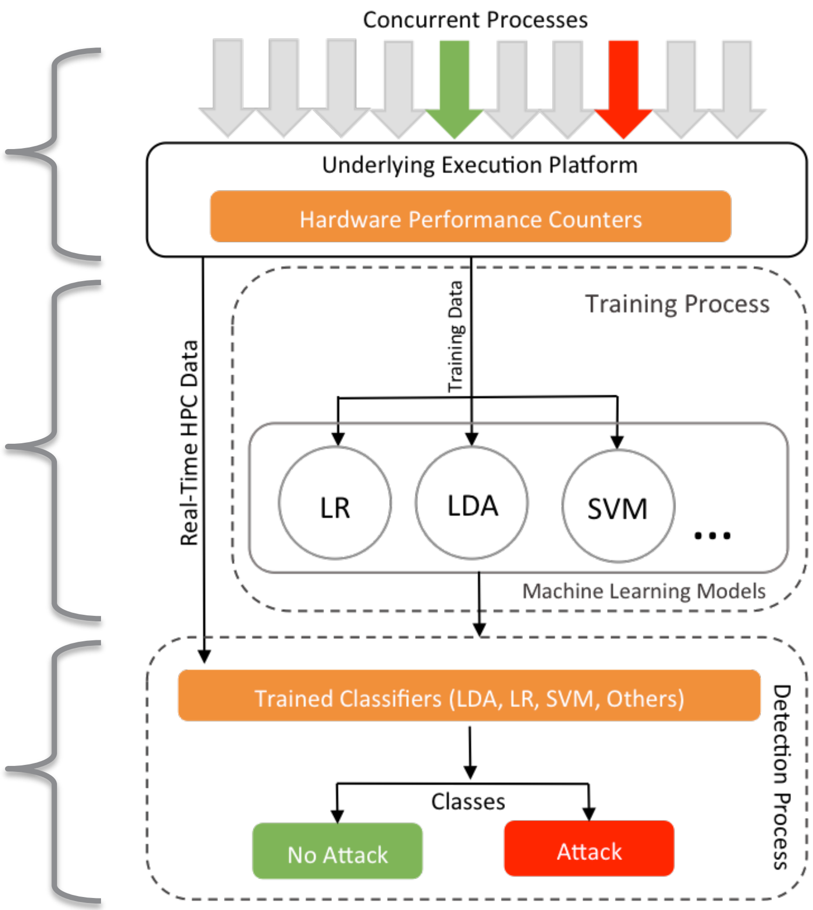

NIGHTs-WATCH

Machine Learning Models

#	Machine Learning Models
1	LR
2	LDA
3	Linear SVM
4	QDA
5	Nearest Centroid
6	Naïve Bayes
7	KNN
8	Perceptron
9	Decision Tree
10	Dummy
11	Random Forest
12	Neural Network

#	Selected Machine Learning Models
1	Linear Regression (LR)
2	Linear Discriminant Analysis (LDA)
3	Linear Support Vector Machine (SVM)

- Linear classifiers could do the job!
- Light-weight for run-time detection
- Easy to embedded with victim process

NIGHTs-WATCH

- ML Models –accuracy for F+R attack detection

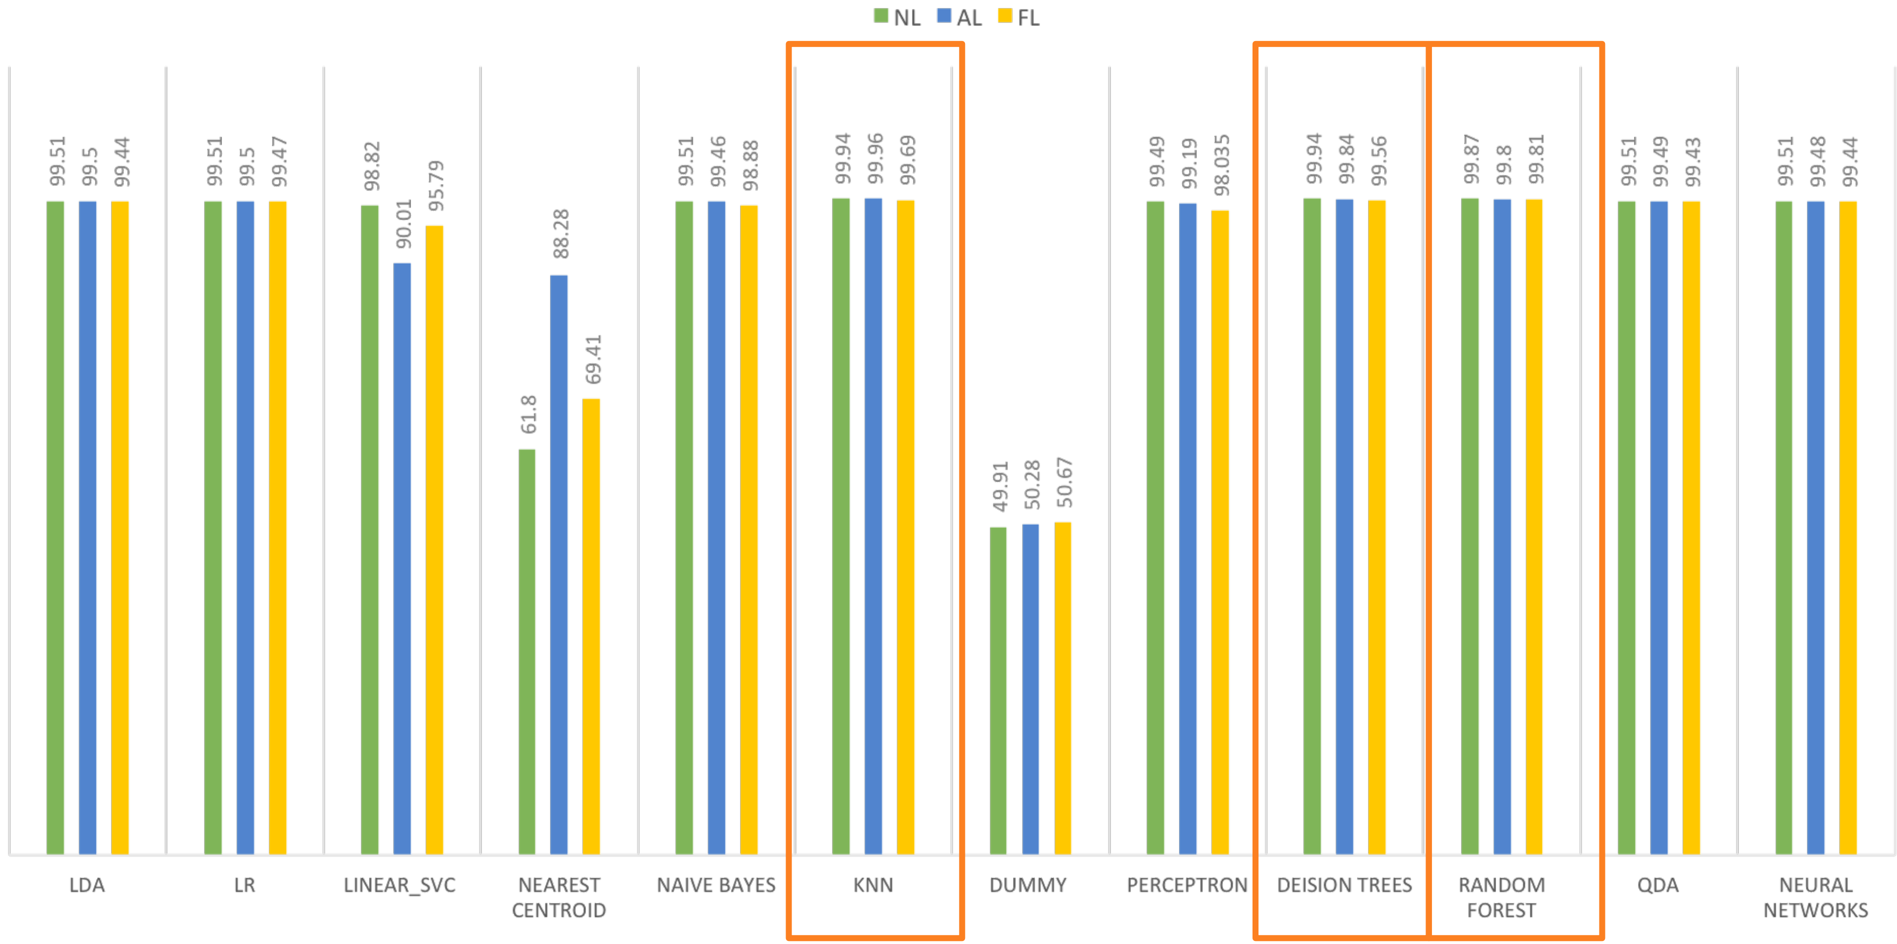

NIGHTs-WATCH

- ML Models –accuracy for F+F attack detection

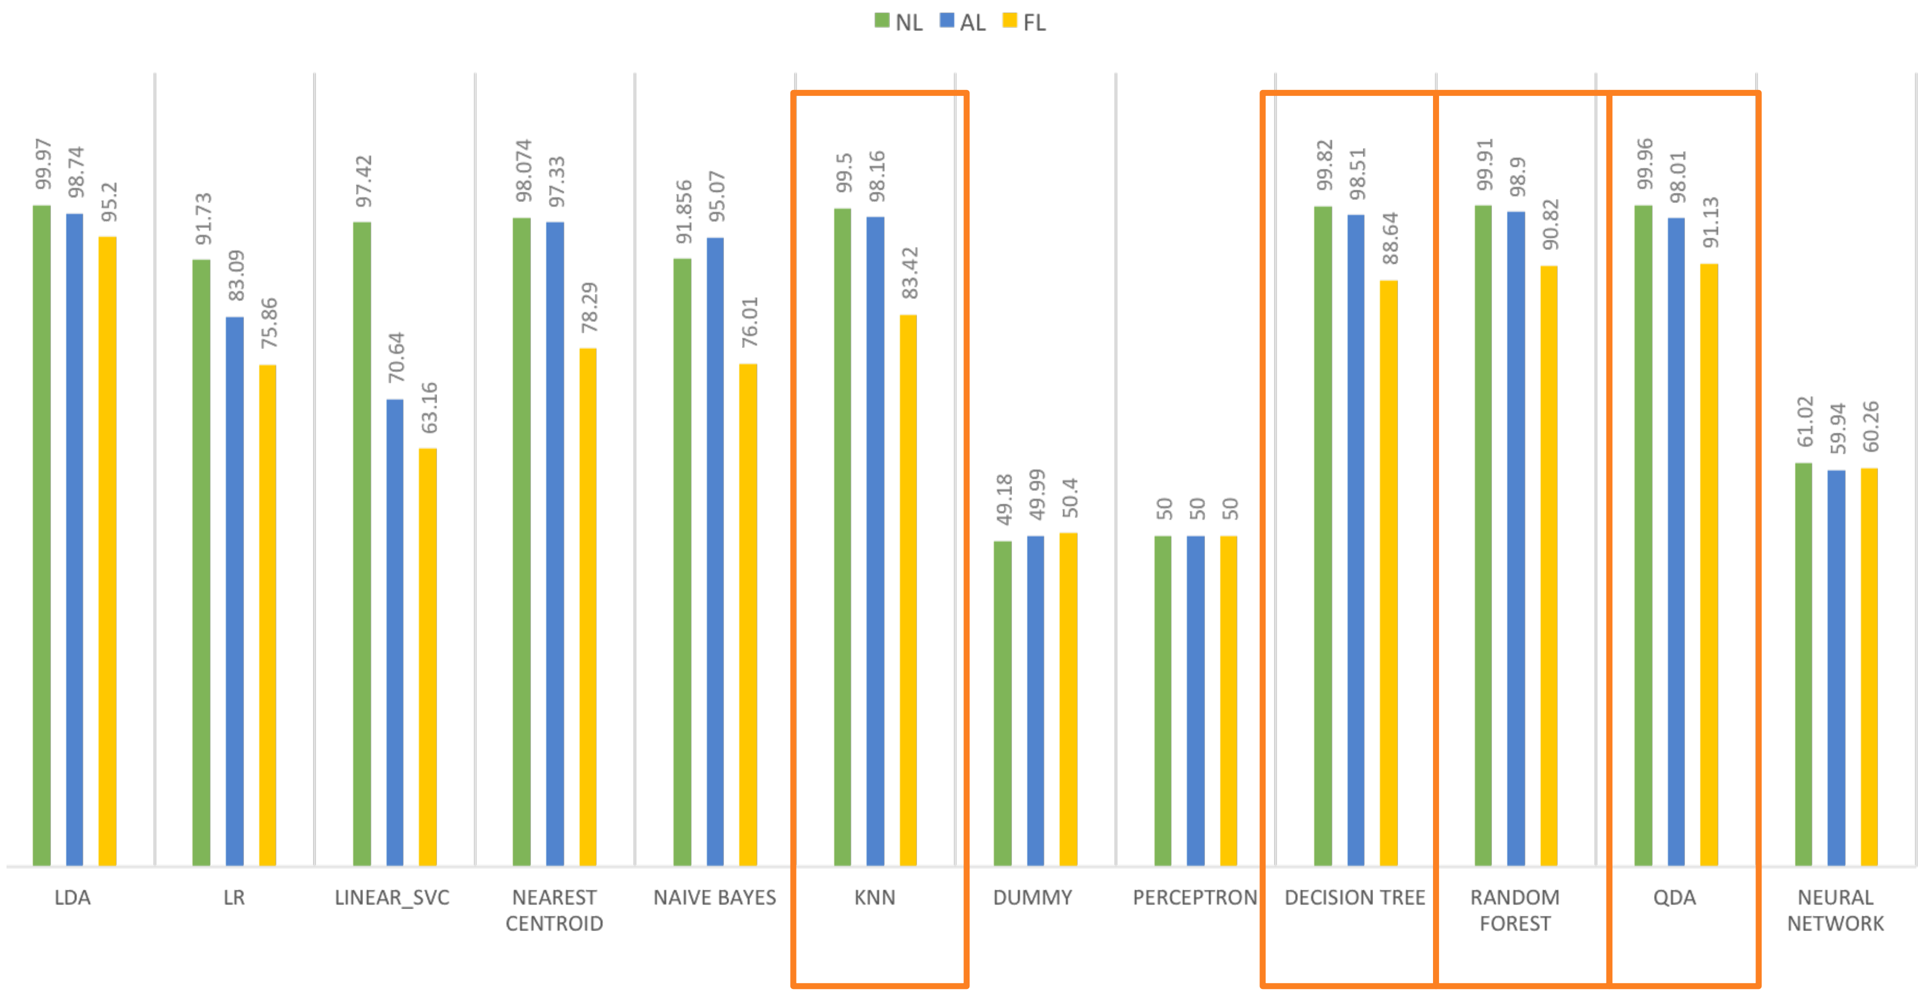

Outline

- Cache-based Side-Channel Attacks
  - The Memory Footprint
  - Hardware Performance Counters (HPCs)
  - Machine Learning Models
- NIGHTs-WATCH
  - The Big Picture
  - Selected HPCs & Machine Learning Models
- Experimental Results
  - Case Study-I: Flush+Reload on RSA
  - Case Study-II: Flush+Flush on AES

Experiments

- The Evaluation Metrics

  ① Detection Accuracy

  ② Runtime Detection Speed

  ③ Runtime Overhead

  ④ System Load Conditions

  ⑤ Distribution of Error (false positives & false negatives)

Experiments

- Case Study-I: F+R Attack on RSA

Experiments

- Case Study-I: F+R Attack on RSA –No Load

ML Model	Accuracy (%)	False Positives (%)	False Negatives (%)
LDA	99.51	99.60	0.40
LR	99.51	100	0
SVM	98.82	33.72	66.28

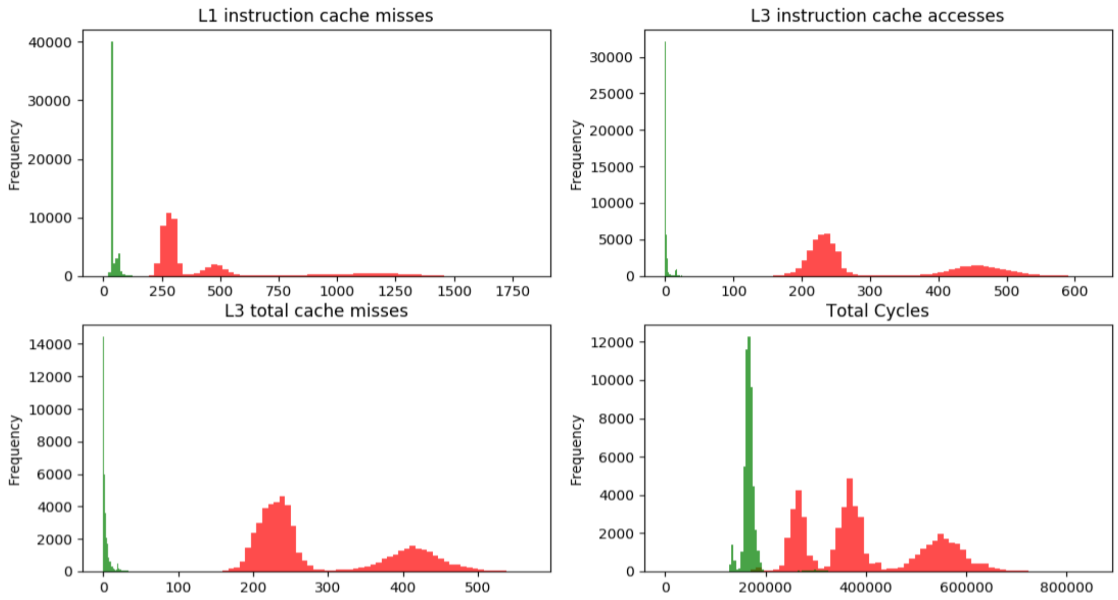

Experiments

- Case Study-I: F+R Attack on RSA –Av. Load

ML Model	Accuracy (%)	False Positives (%)	False Negatives (%)
LDA	99.50	98.42	1.58
LR	99.50	98.82	1.18
SVM	90.01	1.70	98.30

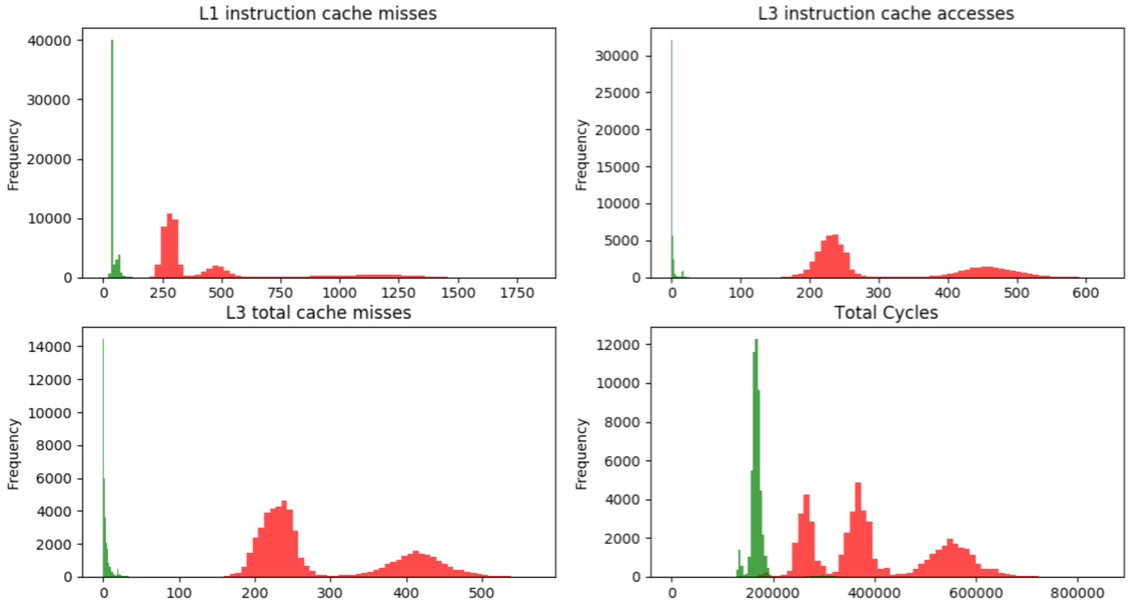

Experiments

- Case Study-I: F+R Attack on RSA –Full Load

ML Model	Accuracy (%)	False Positives (%)	False Negatives (%)
LDA	99.44	87.76	12.24
LR	99.47	92.28	7.72
SVM	95.79	76.29	23.71

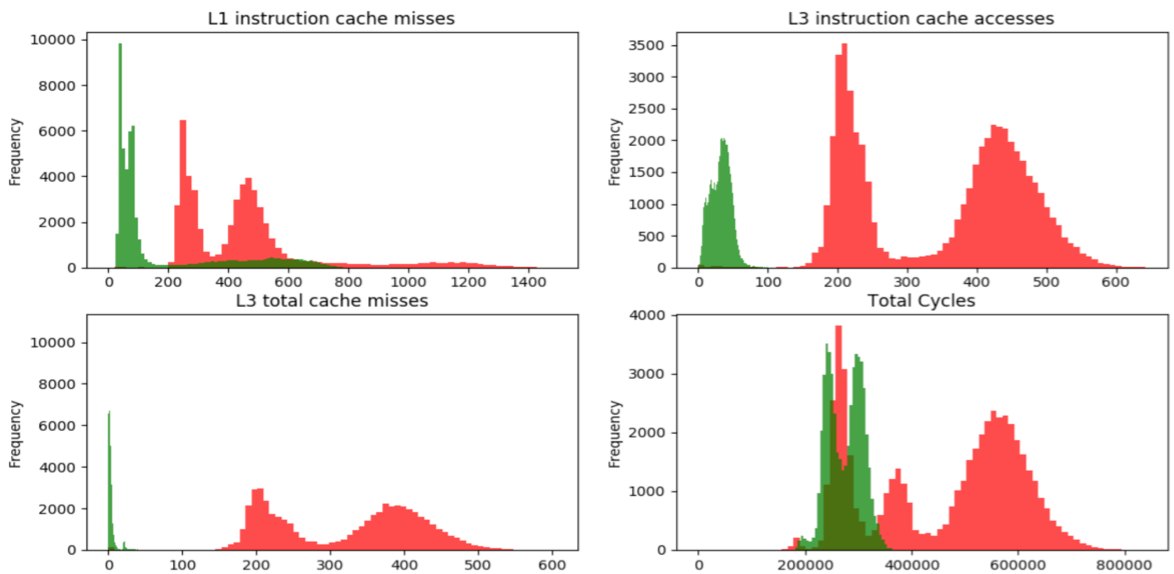

Experiments

- Case Study-I: F+R Attack on RSA
- Speed

ML Model	No/Average/Full Load Conditions
LDA	0.98% of bits are encrypted within single RSA round before successful detection of F+R
LR	
SVM	

- Overhead

ML Model	Victim Slowdown (%)
LDA	0.94%
LR	1.63%
SVM	1.29%

Experiments

- Case Study-II: F+F Attack on AES

Experiments

- Case Study-II: F+F Attack on AES –No Load

ML Model	Accuracy (%)	False Positives (%)	False Negatives (%)
LDA	99.97	75	25
LR	91.73	0	100
SVM	97.42	0	100

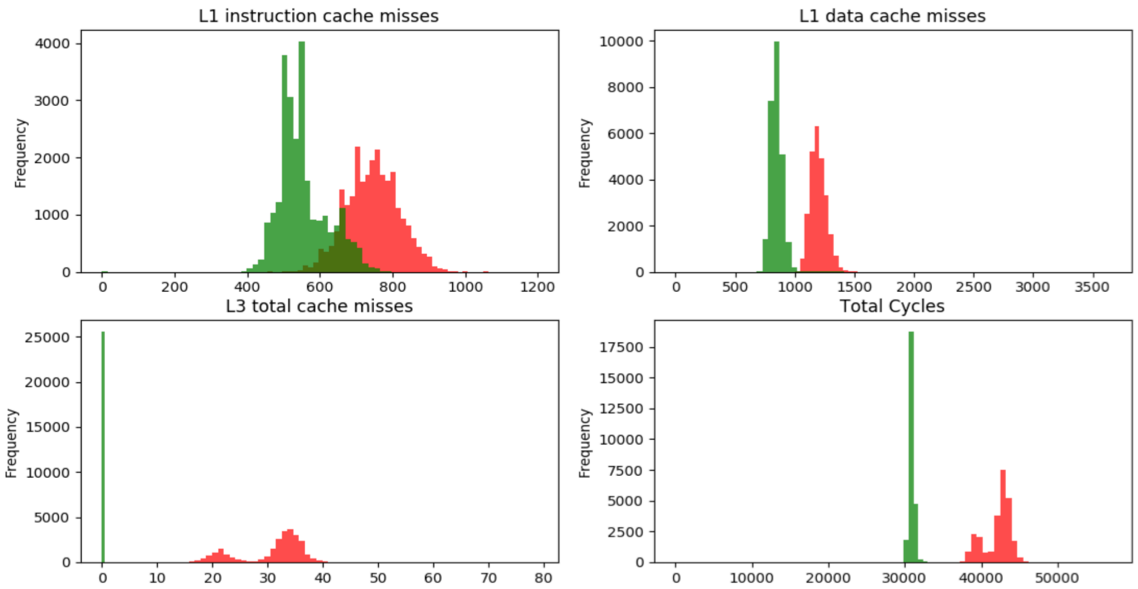

Experiments

- Case Study-II: F+F Attack on AES –Av. Load

ML Model	Accuracy (%)	False Positives (%)	False Negatives (%)
LDA	98.74	89.26	10.74
LR	83.09	84.32	15.68
SVM	70.64	94.56	5.44

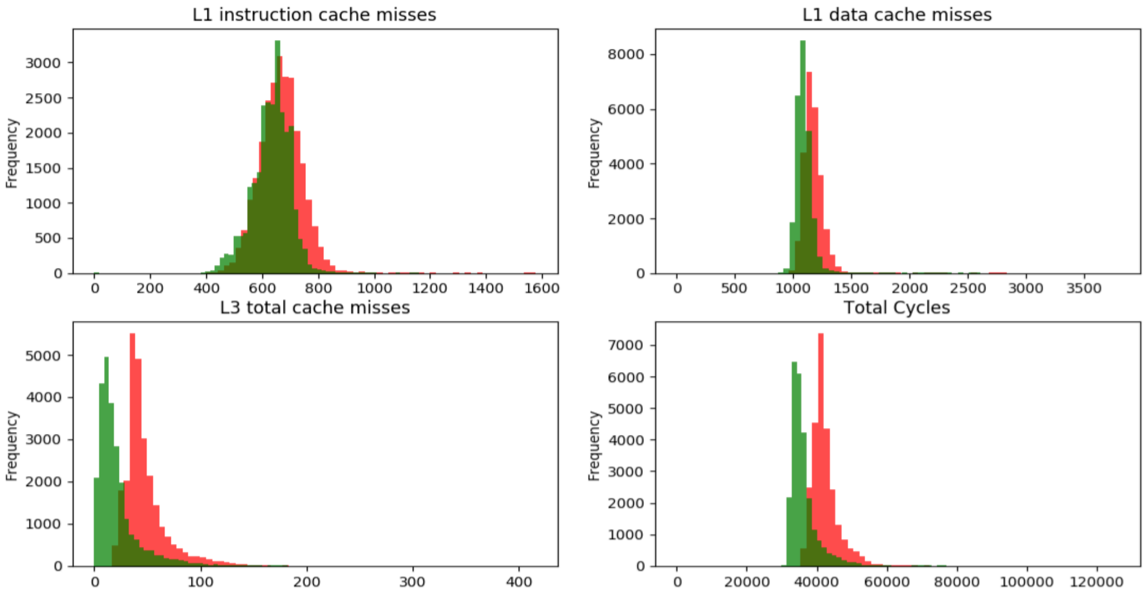

Experiments

- Case Study-II: F+F Attack on AES –Full Load

ML Model	Accuracy (%)	False Positives (%)	False Negatives (%)
LDA	95.20	95.43	4.57
LR	75.86	98.39	1.61
SVM	63.16	98.14	1.86

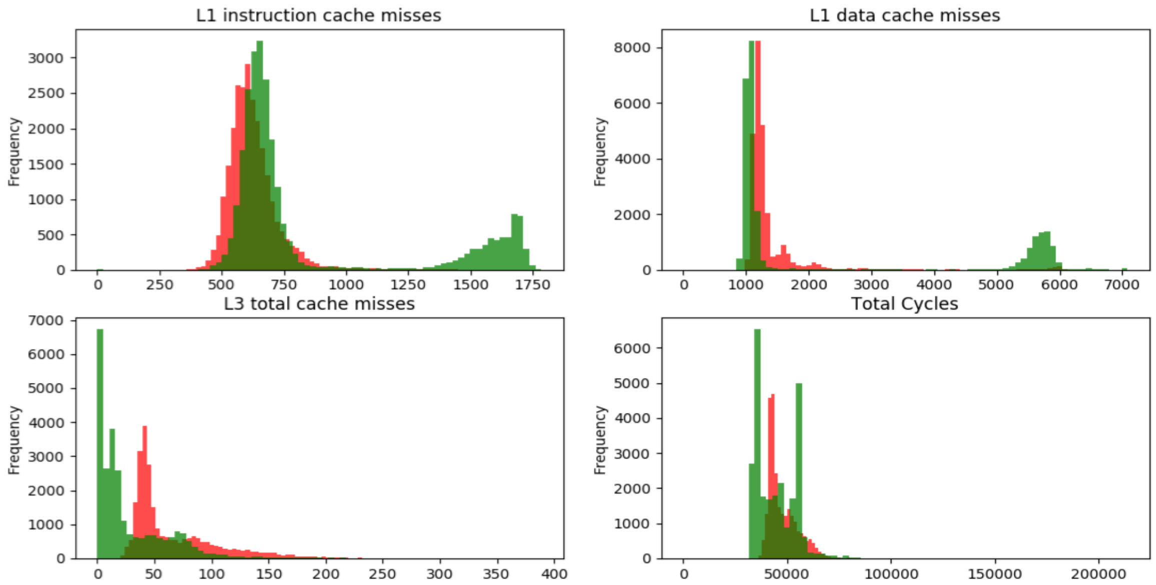

Experiments

- Case Study-II: F+F Attack on AES

- Speed

Technique	Number of encryptions
Flush+Reload	250
Flush+Flush	350
Prime+Probe	4 800

Number of encryptions to determine the upper 4 bits of a key byte.

Gruss et al. 2016. Flush+Flush: A Fast and Stealthy Cache Attack. In DIMVA. 279–299

ML Model	No/Average/Full Load Conditions
LDA	12.5% of 400 AES encryptions are performed before successful detection of F+F
LR	
SVM	

- Overhead

ML Model	Victim Slowdown (%)
LDA	1.18%
LR	1.10%
SVM	0.79%

Concluding Remarks

- NIGHTs-WATCH offers fast runtime detection with high accuracy for cache-based SCAs using machine learning
- Results are consistent under variable load conditions
- Provides detection for high-precision and stealthier attacks on AES & RSA using real-time HPC data
- Scalable for larger set of ML models and attacks

Thank You!

co-authors


Maria (UBS)


Ayaz (WMICH)


Khurram (ITU)


Maham (ITU)


Vianney (UBS)


Guy (UBS)