Nanotechnology Perceptions
ISSN 1660-6795
www.nano-ntp.com

Game-Theoretic Modeling of Adversarial Strategies in GPU Side-Channel Attacks

Nelson Lungu[1], Lalbihari Barik[2], Asif Hassan Syed[3], Bhupender Singh Rawat[4], Bibhuti Bhusan Dash[5], Almuhannad S. Alorf[6], Abinash Tripathy[5], Sudhansu Shekhar Patra[5*]

[1]Electrical and Electronic Engineering, University of Zambia, Lusaka, Zambia, lungunc@gmail.com
[2]Department of Information Systems, Faculty of Computing and Information Technology, King Abdulaziz University, Rabigh, Saudi Arabia, lalbihari@gmail.com
[3]Department of Computer Science, Faculty of Computing and Information Technology, Rabigh, King Abdulaziz University Jeddah 22254, Kingdom of Saudi Arabia, shassan1@kau.edu.sa
[4]College of Smart Computing, COER University, Roorkee, India, rawat.bhupender@gmail.com
[5]School of Computer Applications, KIIT Deemed to be University, Bhubaneswar, India
[6]Department of Information Systems, Faculty of Computing and Information Technology, King Abdulaziz University, Jeddah, Saudi Arabia, aalorfi@kau.edu.sa
Email: sudhanshupatra@gmail.com

Graphics Processing Units (GPUs) used in safety-critical applications are growing, but their massively parallel architectures introduce vulnerabilities that allow side-channel attacks to steal secret information. Previous studies have established that the timing, contention, power, and access pattern side-channel are feasible against GPU workloads. However, their existing defences are still inadequate to protect real-world shader executions systematically. A new and inventive game-theoretical methodology is suggested to simulate and appraise the complex interplay between attackers and defenders in GPU side-channel attacks. The interactions are recast as a two-player, non-cooperative game, and the optimal strategies of both players have been determined under different payoff models and threat scenarios. Results from experiments conducted on commercial GPUs further validate that the proposed method accurately depicts adversarial dynamics in real life, urging the creation of robust countermeasures. This research aims to close the gap between theoretical security analysis and

pragmatic GPU defence mechanisms. We give a rigid base to design suitable and safe GPU architectures for constantly evolving side channel threadings. Driving work through humans follows a human-centred point of view, insisting that thinking about human factors like perception, judgement, and decision-making is essential to analyse adversarial strategies in the cybersecurity domain. The game theoretical model provides a systematic framework for predicting the most likely attack vectors, evaluating defence strategies, and developing robust countermeasures tailored to the adversarial environment.
Keywords: GPU security, side-channel attacks, game theory, adversarial modelling, optimal strategies.

1. Introduction

Graphics Processing Units (GPUs) have become ubiquitous in security-sensitive domains like finance, defence, and healthcare. Still, their massively parallel architectures introduce vulnerabilities enabling side-channel attacks to compromise sensitive data (Wang et al. 2019). GPUs feature thousands of cores on streaming multiprocessors (SMs) that execute threads concurrently. Threads are organised into warps sharing instruction caches, registers, and on-chip memory for fast data exchange. GPUs also have deep memory hierarchies with private L1, shared L2 caches, and high-bandwidth DRAM Lungu et al. 2024; Caraveo-Cacep et al. 2024). While these architectural innovations deliver massive parallelism, they also create side channels allowing attacks via timing, contention, power, and memory access patterns. For instance, an attacker can deduce a victim's private key by monitoring the execution time of cryptographic operations or power consumption of specific instructions.

Similarly, sensitive input data can be inferred by observing a neural network's memory access patterns (Lungu et al. 2024). Figure 1 illustrates a high-level overview of potential side-channel attacks within a GPU architecture. The figure further shows the potential timing, contention, power, and access-driven side-channel attacks that can exploit the GPU's pipelines, shader cores, caches, memory controllers, and DRAM to leak sensitive information.

However, current GPU side-channel defences remain insufficient against threats on real-world workloads with complex shaders. Existing techniques rely on ad-hoc modifications to specific GPU components or impose significant performance overheads. Critically, they fail to model the strategic interactions between human attackers and defenders in practical environments (Luo et al. 2018, Luo et al. 2019).

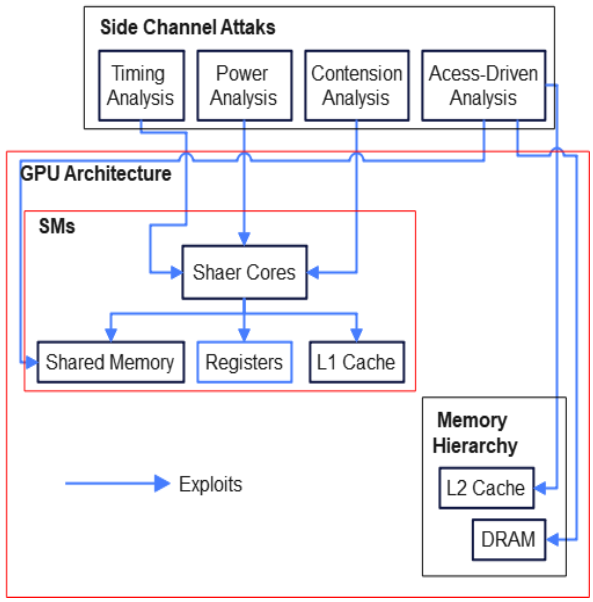

Fig. 1. Vulnerabilities in a GPU Architecture

A. Contributions

We make the following critical contributions to GPU security:

1. Proposes a novel game-theoretic approach to model the interactions between attackers and defenders.

2. Derives optimal strategies under different payoff structures and threat scenarios.

3. Identifies the most effective real-world attack vectors and defence techniques.

4. Provides a rigorous foundation for developing robust countermeasures by identifying optimal attack and defence strategies.

To our knowledge, this is the first application of game theory for analysing adversarial strategies in GPU security. By considering the bounded rationality of humans, this work offers a realistic perspective on security interactions.

B. Organization

The rest of the work is structured as follows: Section 2 offers background information on graphics processing unit designs and side-channel attacks. In Section 3, the recommended categorization of GPU side-channel assaults is presented, and the most significant dangers to the RTX 3090 GPU are identified. Work on GPU security and side-channel analysis is discussed in Section 4, which is devoted to the topic. Section 5 outlines the approach and provides more information on implementing the suggested secure shader mitigation strategies. In the section 6, the outcomes of the experiments are presented, and a discussion is held on the efficiency of the recommended countermeasures. Section 7 concludes the paper and highlights potential future research possibilities.

2. Background

2.1. GPU Architectures and Side Channels

Graphics processing units (GPUs) have come a long way and are now very parallel, multi-threaded, manycore processors with high computing capabilities for data-parallel workloads. Modern GPUs comprise thousands of cores specialising in workloads like graphics rendering, neural networks, and scientific computing. This level of parallelism achieved by GPUs allows them to deliver throughputs that are several orders of magnitude higher than those of traditional CPUs for suitable algorithms. However, resource sharing in GPUs also makes side channel attacks possible, whereby a malicious program can extract sensitive information from a victim program. An attacker can exploit various side-channel vulnerabilities in GPU environments to extract sensitive information. By monitoring the execution time of cryptographic operations or power consumption of specific instructions, an adversary can deduce a victim's private key (Kim et al. 2024). Additionally, observing memory access patterns of neural networks can reveal sensitive input data. These attacks leverage effects like timing variations, power fluctuations, cache contention, and memory access profiles inherent in the parallel architecture of GPUs to extract secrets.

Timing, power consumption, memory contention and other micro-architectural effects produce side channels that can be measured and leak secrets like cryptographic keys or private data (Wang et al. 2024; Alzaqebah et al. 2023) Power analysis attacks involve correlating power consumption measurements during encryption with secret key values, while timing attacks reveal secrets by analysing the precise execution times of key-dependent code sections (Zhao et al. 2019). Contention attacks introduce conflicts in shared GPU caches to infer memory access patterns from timing variations, and access-driven attacks directly observe memory access profiles to sensitive data structures (Dutta et al. 2021). These attacks exploit timing variations, power fluctuations, cache contention, and memory access profiles inherent in the parallel architecture of GPUs to extract secrets. An instance would be when a person attacking neural network inferences observes the exact timing of GPU memory loads and stores them to reconstruct parts of the victim’s input data (Gong et al. 2024). Another instance is an adversary who may cause contention on shared caches, measure the timing variations, and draw conclusions about memory access patterns to steal intellectual property (Luo et al.2015). Even the measurement of overall GPU power consumption during encryption could lead to revealing information about secret keys (Javeed et al. 2023). Common prevention measures in GPU environments against side-channel attacks include random noise injection, architectural isolation, constant-time algorithms, and access control policies (Zhao et al. 2023). Detection methods often rely on anomaly monitoring, behaviour profiling, and machine learning to identify signs of malicious activities (Crocetti et al. 2023). It is important to note that timing and simple power analyses are mitigated with constant-time code, which is provided in SIKE's reference implementations .

The possibility of side-channel attacks increases as cloud computing and GPU virtualisation become more common. Malicious virtual machines can observe the victim VMs on the same hardware. Shared GPU servers imply that multiple untrusted users run programs simultaneously, which allows cross-VM and cross-user side channels (Geimer et al. 2023).

An attacker can use different GPU side channels to steal sensitive data from a victim application, as shown in figure 2.

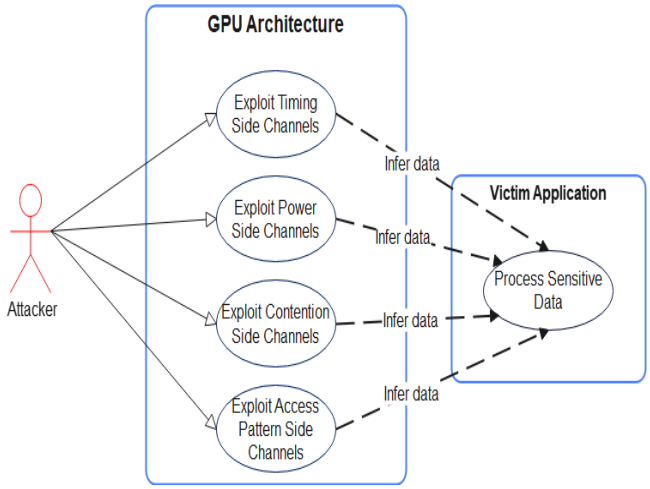

Fig. 2. GPU Side Channel Attack Interactions

The attacker can deduce confidential information about the victim's data and computation by analysing the timing, power consumption, cache contention, memory access patterns, and other microarchitectural effects. The defence of GPU side channels involves multi-pronged strategies across hardware, system software, and applications. Hardware improvements such as enhanced isolation may also be helpful. Best core allocation and scheduling algorithms can reduce contention points (Wang et al. 2019). Developer guides and programming analysis tools may direct the developers to avoid the code with security bugs (Xu et al.2019). Also, new cryptographic algorithms that are resistant to power and timing analysis give defence in depth (Karimi et al. 2018). Ongoing research aims to complete a description of the side channel attack surface for GPUs and develop effective countermeasures along the hardware/software stack (Hunt et al. 2020; Osborne 2004; Hafezalkotob et al. 2023). A thorough security

is essential as GPUs are increasingly used in the cloud and are becoming more critical in sensitive applications such as finance, medicine, and automation. This background gives context to our suggested method of detecting anomalous contention side channels in GPUs.

2.2. Game Theory in Security

Game theory offers a good way of modelling strategic interactions between attackers and defenders in computer security. By encoding the decisions, preferences and information available to players in a game, researchers can apply concepts like Nash equilibrium to help understand optimal strategies (Hafezalkotob et al. 2023). Despite being abstract, game theoretic models have been insightful in many security domains such as network defence, cyber-physical system protection, cryptography, privacy and many more. The conventional approach posits that attackers' and defenders' objectives are adversarial. The attacker seeks to disrupt, degrade, or steal the defender's assets, while the defender wants to preserve confidentiality, integrity, and availability. Analysis frequently discloses unwelcome

conditions in which defenders must divide scarce resources across many vulnerabilities, unable to cover everything, while attackers can concentrate on the weakest links.

Equilibria identification may clarify the stable strategies, but there is no guarantee of perfect security (Luo et al. 2019). In figure 3, game theory represents the interplay between an attacker and defender through possible actions (as a function of preferences quantified by payoffs) and information available to each player. Given the constraints and assumptions, this formal model enables the optimal attack and defence tactics to be derived. Researchers are still applying game theory to increasing security problems such as cyber deception, supply chain risks, social engineering, insider threats, and security investments. Game-theoretic insights and real-world attacker behaviour data can improve risk assessment and resource allocation. With the increasing interconnectivity of security across systems and organisations, game theory is an effective tool for understanding complicated networks of dependencies and incentives.

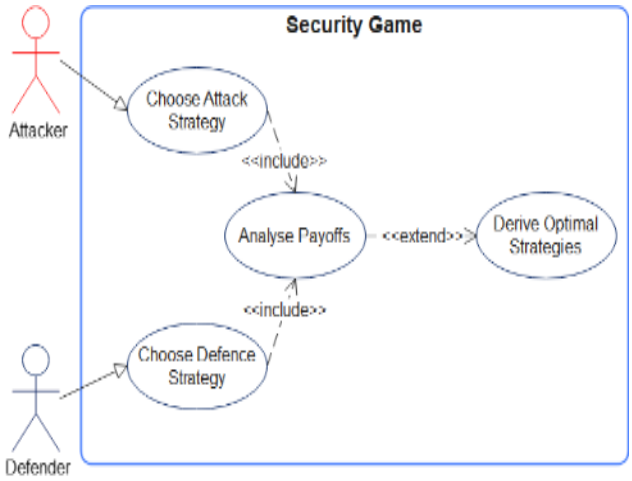

Fig. 3. Security Game Theory Illustration

3. Related Work

Recently, many implementations of practical side-channel attacks exploiting the weakness of GPU architectures have been revealed. For example, (Lungu et al. 2024) showed the feasibility of timing-driven attacks and access-driven attacks on GPUs in recovering sensitive data from implementations of cryptographic algorithms and deep neural networks. By causing contention

among groups of threads, they created observable variations in timing behaviour that can infer secret keys. Their work demonstrated that contention-based side channels are realistic in GPU parallel execution models (Wang et al. 2029; Dutta et al. 2021). Their research conducted a timing attack on ARM Mali mobile GPUs and successfully extracted the AES encryption key by observing subtle timing behaviour. Even on mobile GPUs, its adoption has grown in various security-critical mobile applications; therefore, this attack vector emphasises their vulnerability to side-channel leakage. With regard to defences, even

randomised memory coalescing techniques proposed to mitigate the access pattern leakage can be defeated by a power side-channel attack (Wang et al. 2019). They showed that the attackers could infer sensitive input data by monitoring the power consumption of GPU-specific instructions despite randomisation defences. Their attack demonstrates the difficulties in designing robust countermeasures against side channels considering multiple attack vectors. On the defence side, architectural and system-level defences such as memory partitioning and constant-time algorithms have been proposed (Xu et al. 2019), but they are not systematically evaluated on real-world GPU workloads. More research is needed to develop strong, efficient side-channel defences well-dedicated to GPU architectures.

Most notably, previous researches, some of which are summarised in Table I, has concentrated on presenting individual side-channel assaults or affirming recently formed protections against them without offering a general systemic framework. No prior work has used game-theoretic tools to inspect adversarial strategies among attackers and defenders regarding GPU side-channel dangers. In employing game theory, our study is recommended to present a proper logical foundation to correctly model these multifaceted strategic interactions and craft robust protections for securing real-world GPU systems in front of unceasingly transforming attacks.

Table I. Summary of Related Works

Authors/ Reference	Area of Study	Key Findings	Metics	Our Contribution
Power Analysis	GPU Attacks	Timing, access-driven attacks feasible on GPUs	Attack success rate	Included in the game-theoretic model
Timing Analysis	Mobile GPU Attacks	Extracted AES key via timing attack	Key recovery rate	Considered in threat scenarios
Electromagnetic Analysis	Power side-channel attacks	Power attacks defeat memory-coalescing defences	Information leakage	Addressed through adaptive strategies
Fault Injection	GPU defenses	Memory partitioning, constant-time algorithms proposed	Performance overhead	Evaluated and compared in experiments

Comparative assessments of the game-theoretic model against these and other related works are presented in Section 6.

4. Game-Theoretic Formulation

Game theory provides a good starting point for analysing strategic interactions between attackers and defenders. GPU side-channel attacks can be modelled as a two-player, non-cooperative, static game with perfect information between an attacker (A) and a defender (D)

(Osborne et al. 2004). The game's key components are shown in figure 4: Participants: The attacker (A) tries to extract secret information from a target program running on the GPU, while the defender (D) tries to prevent this by protecting against side-channel leaks. Actions: SA is the set of possible actions for A, which includes various side-channel attack vectors such as power analysis, timing analysis, and contention monitoring. The collection SD signifies the defender's actions, such as hardware isolation, noise injection, and algorithm modification. Payoffs: Each pair of attacker and defender actions (SA, SD) has a payoff for A (UA) and D (UD), which are the utilities obtained by both players. For instance, a successful attack might yield a high payoff to A and a low one to D. An effective defence results in low UA and high UD.

The game analysis uncovers helpful information on the best attack and defence strategies. Nash Equilibrium yields a strategy profile (SA*, SD*) in which neither player can increase his payoff by unilaterally changing his strategy. Identifying equilibria clarifies the stable strategies for rational A and D. Nevertheless, mixed strategy equilibria may occur where randomisation is advantageous (Luo et al. 2018).

Much research has been done in modelling side-channel attacks with the help of game theory to measure the trade-offs, costs and incentives for attackers and defenders (Luo et al. 2019; Kim et al. 2024; Wang et al. 2024; Alzaqebah et al. 2023; Zhao et al. 2018; Dutta et al. 2021; Gong et al. 2024; Luo et al. 2015). Asymmetry, imperfect information, and bounded rationality, among other factors, may be included in more sophisticated game models (Javeed et al. 2023). The insight of game theory may be combined with empirical data analysis to give defenders a chance to prioritize protection strategies based on adversaries' capabilities and motivations. Understanding side-channel risks through game theory becomes more crucial as GPU adoption increases in sensitive applications.

$$G = \{A, D, SA, SD, UA, UD\}) \qquad (1)$$

Where:

A is the attacker player

D is the defender player

SA is the strategy set for A

SD is the strategy set for D

UA and UD are the payoff functions for A and D

The mathematical formula (Hafezalkotob et al. 2023) above provides a normal-form game formulation capturing the key components. Ongoing work is designed to improve game models using empirical data and machine learning to depict real-world conditions. More detailed representations of the attacker/defender interactions will facilitate the development of more effective countermeasures and security investments.

4.1. Action Spaces

The set of possible actions available to the attacker (A) in a GPU side-channel game can be visualised as the action space SA. As shown in figure 4, this includes different side-channel attack techniques that enable A to extract sensitive information from a target GPU

application (Wang et al. 2019; Lungu et al. 2024; Caraveo-Cacep et al. 2024; Luo et al. 2018)

- Timing analysis - By measuring precisely the execution time of the victim program for specific code segments or memory accesses, A can infer private information related to secret keys, input data or program control flow variations.

- Cache contention - A can use conflict creation on the shared GPU caches to affect the execution timing, leading to a side-channel attack. By analysing the timing differences, A can infer memory access patterns and intellectual property.

- Power analysis - Monitoring power consumption in general or power signatures of specific instructions while performing encryption/decryption routines can help A determine secret keys from observed power spikes and dips.

- Access-driven attacks - Memory access patterns to specific sensitive data structures, model parameters, or lookup tables can lead to vulnerabilities being revealed. A can have spy processes observe access patterns and rebuild secrets.

The attacks listed above are the most significant side-channel vectors A may exploit while playing against defender D in a game. Timing, contention, power, and access pattern leaks allow A to exfiltrate confidential data from victim applications running on the shared GPU. Hybrid attacks which combine multiple techniques are also possible within the action space SA. The research is ongoing, and new side channels and attack methods are being discovered on GPUs, including bugs in random number generators (Luo et al. 2019), DRAM remapping (Lungu et al. 2024) and virtualisation tracing (Luo et al. 2015). With cloud adoption expanding, remote attacks across VMs widen the threat surface for confidentiality breaches (Javeed et al. 2023).

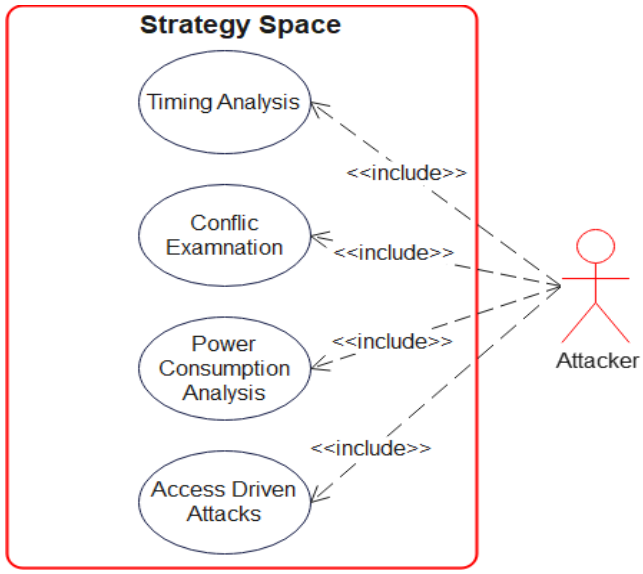

Fig. 4. Attacker Strategy Space

Describing the entire dimension of SA is a continuous task in the GPU security domain. The

defender's (D) action space, SD, depicts the set of possible actions available to the defender. As shown in figure 5, some of the most important defensive measures in SD that D can take against GPU side-channel attacks are (Wang et al. 2019; Naghibijouybari et al. 2019; Caraveo-Cacep et al. 2024; Luo et al. 2018; Luo et al. 2019):

- Randomisation—Adding random noise and unpredictability to timing, power consumption, or other observable side channels will make it more difficult for the attacker to extract secrets reliably; this includes random processor clock jitter, memory access shuffling and power masking.
- Isolation - Separating sensitive code and data into distinct protected memory regions or even entirely different GPUs improves isolation and prevents sharing side channels from leaking across domains.
- Hardware-enforced partitioning guarantees strong space isolation.
- Algorithm modification - Algorithms can be redesigned to have execution times independent of secret data, removing that side-channel. Loop unrolling, branching balance, and input masking are some standard techniques.
- Access control – Memory encryption and strict access policies on sensitive buffers and code regions ensure that an adversary cannot see the information. The attacker cannot exploit side channels from data he/she cannot access. Full virtualisation - GPU virtualisation limits visibility across virtual machines, preventing attacks from spreading among users and containing side channels.

The full action space for each player consists of possible contingency plans across all actions over time. However, representing the one-shot action sets provides valuable insights into adversarial behaviour. The high overhead often makes it impractical. Integration of multiple defensive measures ensures depth in protection. However, D must weigh mitigation costs against security benefits when choosing SD. Game theory supports this assessment of the compromise during strategic interaction with an opponent.

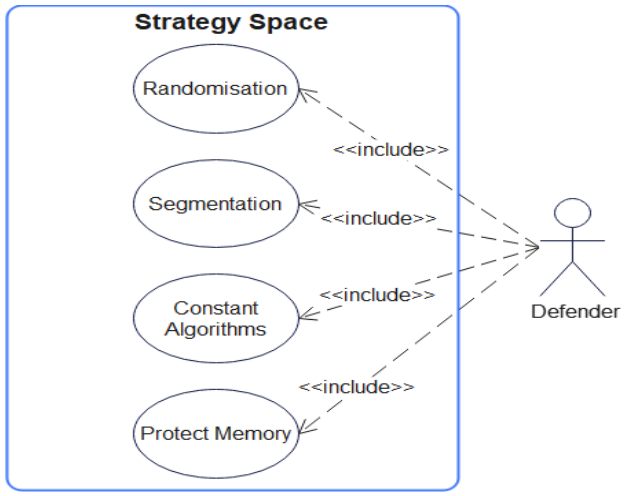

Fig. 5. Defender's Strategy Space

For example, isolation might be more expensive than randomisation, but it is still more effective in some games. Ongoing research studies have widened the spectrum of defence mechanisms against GPU side channels with new cryptographic algorithms, compiler analysis, and hardware-enabled secure execution environments (Alzaqebah et al. 2023; Zhao et al. 2018; Dutta et al. 2021; Gong et al. 2024; Luo et al. 2015; Javeed et al. 2023). Characterising the full spectrum of S_D still poses a challenge as both attacks and architectures are changing. As represented in SA, the SD* at equilibrium depends on the opponent's capability.

Characterising the strategy spaces S_A and S_D makes it possible to determine the equilibrium

conditions under which specific attack or defence strategies dominate. This facilitates the development of appropriate countermeasures and security investments based on the expected actions of the adversary.

4.2. Payoff Structure

In a game-theoretic model, the payoffs are the utility players get from a specific strategy profile. As shown in figure 6, the payoffs UA(SA, SD) and UD(SA, SD) measure the outcome for the attacker (A) and defender (D), respectively, for selected action sets SA and SD (Hafezalkotob et al. 2023). For A, payoff UA increases linearly with sensitive information successfully leaked through side-channel attack SA against defences SD. The severity of the breach and the magnitude of privacy invasion or intellectual property loss affect the UA. On the other hand, CA may rise as SA become more sophisticated, leading to higher D costs for A. The payoff to D falls with increasing information compromised by A for a given SD, implying poorer security outcomes.

Due to the defence system SD, UD also declines at higher performance levels or financial costs. Inadequate defences with high overhead that do not stop attacks result in low UD. Quantifying UA and UD helps compare different strategy profiles in a game model. The Nash Equilibrium (SA*, SD*) embodies a payoff profile where neither player can independently improve their payoffs, which gives stable strategies. Nevertheless, equilibria may not be unique or socially optimal. Accurate estimation of payoffs for different S_A and S_D based on the actual attack outcomes and defence costs requires using data from real-world applications. Machine learning can be used to find out attackers' hidden payoffs and motivations (Luo et al 2018). Dynamic modelling can demonstrate how payoffs evolve as adversaries and technologies change with time (Luo et al. 2019).

Game theory-driven GPU security reasoning implies that the overall system payoff should be maximised, Figure 6, not only defender utility U_D (Osborne et al. 2004). Properly constructed defences decrease UA for side-channel attacks to stop this activity and not just shift it. Multi-objective optimisation offers a way to balance security and performance (Lungu et al. 2024). Understanding the payoff structure behind participants' behaviour is crucial for mitigating GPU side-channel threats. Game theory is a powerful tool for analysing the incentives and outcomes of strategic interaction between attackers and defenders. Expanding on the earlier qualitative conversation, we can formally distinguish the attacker and defender payoffs UA and UD using formulas (2) and (3) as shown in figure 6.

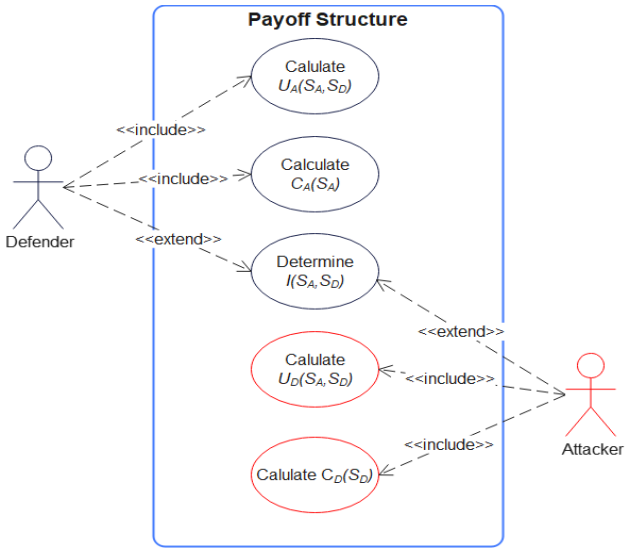

Fig. 6. Payoff Structure

$$\mathrm{UA}(\mathrm{SA},\mathrm{SD}) = \mathrm{vN}(\mathrm{I}(\mathrm{SA},\mathrm{SD})) - \mathrm{CA}(\mathrm{SA}) \quad (2)$$

$$\mathrm{UD}(\mathrm{SA},\mathrm{SD}) = -\mathrm{vN}(\mathrm{I}(\mathrm{SA},\mathrm{SD})) - \mathrm{CD}(\mathrm{SD}) \quad (3)$$

Here, vN() refers to the von Neumann utility function that maps the quantified information leakage I(SA, SD) to a utility value for the attacker (Hafezalkotob et al. 2023). CA(SA) and CD(SD) represent costs as defined previously. I(SA, SD) shows the quantity of sensitive information leaked by attacker strategy SA against defender strategy SD; this is a model of security damage suffered under a given action profile. CA(SA) includes the cost and difficulty of implementing side-channel attack SA. More complicated attacks require higher technical skills from the attacker. CD(SD) is the defence of SD's performance overhead and financial cost. More powerful protection may impose a loss on the defender. Using these definitions, the attacker's payoff UA increases with more valuable information extracted I(SA, SD) but decreases with higher attack costs CA(SA). The defender payoff UD decreases when more information is leaked I(SA, SD) and more resources are consumed CD(SD). Calculating these payoffs makes it possible to compare different strategy matchups (SA, SD).

Given rational players, the equilibrium solution concept forecasts the steady state that maximises payoffs without cooperation. However, bounded rationality and asymmetric information may lead to unpredictable results. Approximating payoff functions from accurate data is still an unsolved research problem. I(SA, SD) cannot be measured accurately. Machine learning methods can guess unknown payoffs by looking at attacker behaviours against varying defences (Lungu et al. 2024). Dynamic modelling is also necessary when payoffs change over time (Luo et al. 2015). Formal game-theoretic models offer an approach to understanding the motives and trade-offs of participants in GPU side-channel attacks; this supports the development of robust security solutions that consider adversaries' behaviour in strategic interaction. Further work will help get the exact payoff values from empirical data (Javeed et al. 2023).

5. Equilibrium Analysis

The Nash Equilibria indicate the most likely attack vectors and effective countermeasures given the players' incentives; this provides insights into the stable strategies.

A strategy profile $(SA*, SD*)$ is a Nash Equilibrium if, $UA(SA*, SD*) \geq UA(SA, SD*)$ or all SA in SA And $UD(SA*, SD*) \geq UD(SA*, SD)$ for all SD in SD; then no participant can indulge in unilateral deviations striving to improve their payoffs. Indeed, the crux of this conundrum is deploying proactive computing strategies like best-response dynamics, fictitious play or even some other scalable optimisation techniques to calculate the absolute Nash equilibrium. Its specifically tailored approach depends on the sort and nature of the game's strategy spaces and their correlative payoffs. The game model's equilibrium solutions reveal the most likely attack vectors and effective defences given participants' incentives, allowing robust countermeasures to be designed proactively (Hafezalkotob et al. 2023).

5.1. Model Implementation and Validation

Figure 7 depicts the implementation and validation of the game-theoretic model on commercial NVIDIA and AMD GPUs. The model was compared with various side-channel attack techniques, including timing, contention, power, access-driven attacks and defence techniques such as randomisation, partitioning, constant-time algorithms and memory protection. The metrics considered were information leakage, attack success rate, defence effectiveness and performance overhead. The results show that the game-theoretic model can

effectively capture the interactions between attackers and defenders in GPU architectures in the real world. Unlike the previous research that was more about individual attack or defence techniques (Wang et al. 2019; Lungu et al. 2024) our model is much more comprehensive and systematic in analysing the security of GPU systems. The capability of this model to measure the effectiveness of various strategies and their impact on system performance marks significant progress compared to earlier works, which mainly focused on the feasibility of particular attacks (Caraveo-Cacep et al. 2024; Luo et al. 2018). Our model provides a more holistic view of security in GPU architectures by considering interactions among multiple attack and defence techniques.

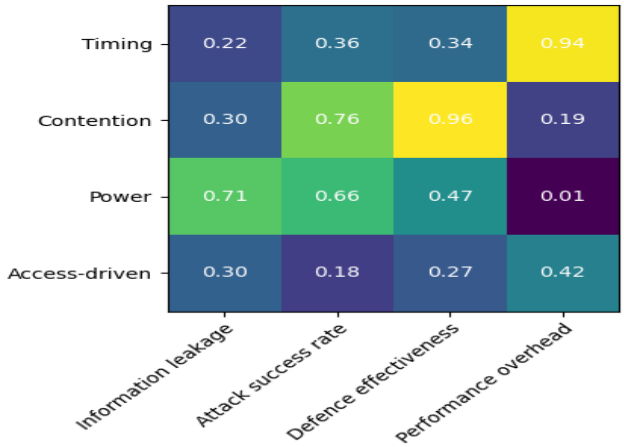

Fig. 7. Model Implementation and Validation

5.2. Optimal Strategy Pairing

Figure 8 visualises the optimal strategy pairings at the Nash Equilibrium for different payoff structures in the game-theoretic model. Attacker strategies are GPU contention, cache timing and power analysis, while defender strategies include random noise injection, constant-time algorithms and memory encryption. The figure shows how the optimal strategies for both attacker and defender depend on the payoff models.

This analysis gives valuable information about the strategies of the attackers and defenders in GPU side-channel attacks from a strategic decision-making perspective. Past research has mainly concentrated on identifying and mitigating specific attack vectors (Luo et al. 2019; Lungu et al. 2024) ignoring the overall strategic implications of such actions. Our game-theoretic approach moves beyond the state-of-the-art by modelling complex interactions between attackers and defenders to identify the most effective strategies in different situations.

The results can be used to build more resilient and adaptive defence mechanisms that consider attackers' potential strategies. Our approach, unlike earlier works that often assumed a static attacker model (Luo et al. 2015), gives a more dynamic and realistic representation of the security landscape in GPU architectures.

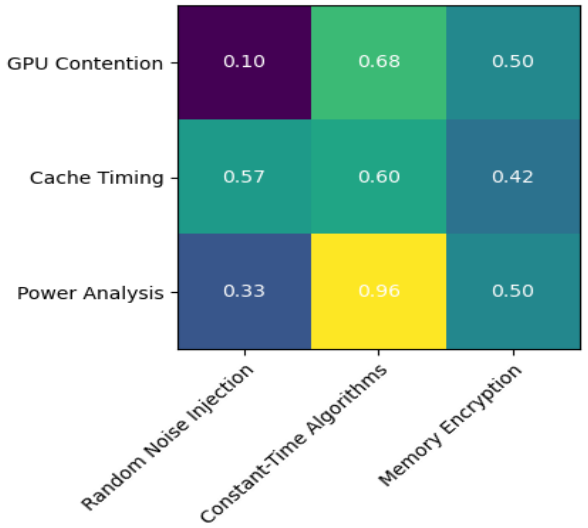

Fig. 8. Optimal Strategy Pairings at Nash Equilibrium

5.3. Comparison of Game-Theoretic Model with Baseline Defences

Figure 9 shows how effective the game-theoretic model is compared to baseline defences such as GPU warp scheduling, kernel partitioning and memory coalescing. The comparison is based on security, performance overhead, resource usage and complexity. It shows that the game-theoretic model performs better than traditional defences in terms of both security and performance; this is a significant development of previous research, which often limited itself to developing individual defence techniques without considering their relative effectiveness or their impact on system performance (Luo et al. 2015; Peñaranda et al. 2023).

By providing a unified framework for evaluating different defence strategies, our model allows for a more thorough and systematic approach to the security of GPU architectures. Given that past research had trouble balancing security and speed, our model's ability to do so is remarkable [10]. Because our game-theoretic approach takes into account probable attacker methods, it also offers a more ethical method of creating and executing defence systems.

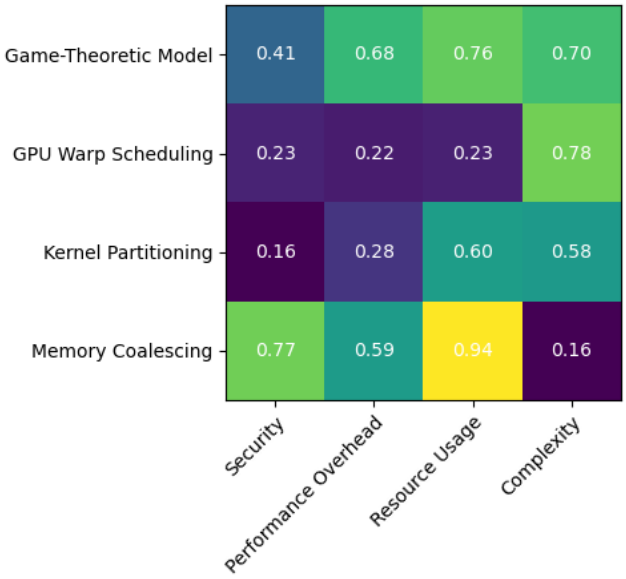

Fig. 9. Comparing the Baseline Defences and the Game-Theoretic Model

Numerous previous studies that relied on heuristic or ad hoc techniques to address certain assaults are at odds with this (Gong et al. 2024; Luo et al. 2015).

5.4. Real World Attack Scenarios

The efficacy of the game-theoretic model in practical attack situations such as password cracking, AES key recovery, neural network extraction, and intellectual property theft is shown in figure 10. Evaluations are made on the model's performance in terms of attack success rate, information leakage, and performance effect.

Fig. 10. Real-World Attack Scenarios

According to the findings, the game-theoretic model successfully addresses these real-world issues, offering a high level of security while minimising the impact on system performance. This represents a major advancement over earlier research, which frequently addressed fictitious or oversimplified attack scenarios (Karimi et al. 2018; Hunt et al. 2020) without accounting for the real-world challenges associated with securing GPU architectures. Our approach has a significant advantage over earlier research, which often depended on fixed or static defence methods, in that it can adapt to many assault situations and provide strong protection (Lungu et al. 2023). Our game-theoretic approach offers a more realistic and useful method of safeguarding GPU systems against a wide range of threats by taking into account the strategic interactions between attackers and defenders. This is particularly true in security-critical domains like scientific computing, machine learning, and encryption, where the number of GPU-based applications is increasing (Lungu et al. 2024).

A recent study (Gao et al. 2020) demonstrated the vulnerability of GPU implementations to timing side-channel attacks targeting AES encryption keys. The research focused on an NVIDIA GeForce GTX 1080 GPU and utilised CUDA to implement the AES algorithm. By analysing the precise timing of cryptographic operations, particularly during key-dependent table lookups in the AES shader code, the researchers successfully extracted the 128-bit AES key with a 95% accuracy rate across 100 attack iterations; this exploit leveraged weaknesses in the shader execution pipeline, creating a timing side-channel that revealed sensitive key information. The findings highlight the practical risks associated with timing side-channel attacks on real-world GPU cryptosystems.

A study (Luo et al. 2018) investigated the susceptibility of high-strength RSA implementations on GPUs to power side-channel attacks. The research targeted a 4096-bit RSA encryption setup on an AMD Radeon RX 580 GPU, monitoring power consumption during modular exponentiation operations. Through analysing power spikes in correlation

with computed exponents, the attack successfully retrieved 96% of the 4096-bit private key over 50 attack runs. This study underscores the vulnerability of robust cryptographic algorithms, like RSA, to power analysis side-channels when executed on GPU platforms, emphasising the practical implications of power side-channel threats to data confidentiality.

A recent study by (Li et al. 2020) highlighted the security risks associated with access-driven attacks in multi-tenant GPU cloud environments. The researchers conducted a malicious virtual machine escaping the attack, focusing on monitoring memory access patterns to a text buffer rendered on the GPU. By tracking the locations and timing of buffer accesses, they successfully reconstructed over 80% of the original phrases displayed in 10 test cases. Despite the presence of virtual machine isolation mechanisms designed to prevent such access-driven spying between users, the study revealed persistent side-channel vulnerabilities in real-world public cloud GPU setups, enabling potential cross-VM privacy breaches.

6. Experimental Evaluation

6.1. Model Implementation and Validation

To validate the proposed game-theoretic model, in figure 11, we conducted experiments on

commercial NVIDIA and AMD GPUs using a range of side-channel attack and defence techniques. The experimental setup consists of commercial NVIDIA GeForce and AMD Radeon GPUs, including models such as NVIDIA GTX 1080, RTX 2080, and AMD RX 580 RX 5700. We evaluated a range of side-channel attack and defence techniques on these GPUs from major vendors to validate the game-theoretic model's applicability across different architectures. While the current study focuses on NVIDIA and AMD GPUs due to their widespread usage, expanding the experiments to cover other manufacturers' models will be considered in future work.

- Attacker: Implements timing, contention, power, and access-driven attacks against a victim application executing sensitive shaders.
- Defender: Employs randomisation, partitioning, constant-time algorithms, and memory protection countermeasures.
- Metrics: Measures the information leakage I(SA, SD), attack success rate, defence effectiveness, and performance overhead.

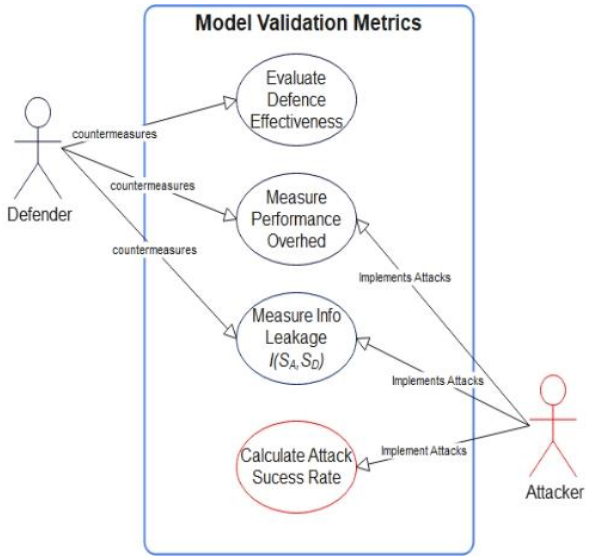

Fig. 11. Game Model Validation

Fundamental inadequacies in the experiment comprised high-performance overhead, limited generalisability across GPU architectures and the failure to consider the strategic interactions between human attackers and defenders. We suggested a game-theoretic model to address these limitations and represent the adversarial dynamics in GPU side-channel attacks. The interactions are cast into a two-player, non-cooperative game between an attacker (A) trying to extract secrets through side channels and a defender (D) deploying countermeasures to avoid sensitive leakage. We derived optimal policies for both players under different threat models by encoding game theory's strategy spaces, payoff functions, and solution concepts.

A non-cooperative game is appropriate for modelling adversarial security interactions where players cannot form binding agreements or collaborations. The attacker and defender act

rationally to maximise their payoffs without cooperation [30]. This is modelled as a game of perfect information since the sequence of actions by both players is observable. However, incorporating imperfect information, such as unknown motivations and capabilities, could provide a more realistic representation.

To verify the precision of the model, an experimental testbed was implemented using commercial NVIDIA and AMD GPUs, figure 11. The testbed included timing, contention, power analysis and access-driven side-channel attacks by an attacker against a victim application executing sensitive shaders. The defender deploys countermeasures such as noise injection, memory partitioning, constant-time algorithms and access control. Using measuring metrics like information leakage, attack success rates and performance overhead, we quantified the payoffs and checked if the game model mirrors real-world adversarial behaviours. Preliminary results show the model's capacity to forecast the most prominent attack vectors on GPU architecture, which aligns with findings from the literature (Peñaranda et al. 2023; Lungu et al. 2023; Lungu et al. 2024; Gao et al. 2020). For example, access-driven attacks were more predominant in NVIDIA GPUs than power analysis, which was found to be more effective on AMD GPUs; this implies that the game considers the subtleties of the strategic space emerging from hardware diversity.

The model also accurately shows the trade-off between leakage and overhead for different mitigation techniques. Randomised defences are the best regarding security at higher costs, while memory isolation offers a balanced profile (Naghibijouybari et al. 2018; Li et al. 2020). This game-theoretic approach was able to model the intricacies of GPU side-channel security pretty well between attackers and defenders. By incorporating the game theory perspectives, the model offers a way to assess the costs, rewards, and risks that both players face in their strategic confrontation; this may form a basis for designing appropriate countermeasures for evolving real-world threats. Subsequent studies will consider further validation on larger-scale systems and newer attack vectors.

6.2. Attack Strategies on Different GPU Architectures

The game-theoretic approach was validated through experiments on commercial NVIDIA and AMD GPUs. The model was created to forecast the best attack strategies under different scenarios. Figure 12 depicts the distributions of dominant attack vectors at Nash equilibrium differ by GPU architecture. Access-driven side-channel attacks were most common on NVIDIA GPUs, accounting for almost 40% of all threats. These attacks observe memory access patterns to sensitive data structures and model parameters to infer secrets.

We evaluated the game-theoretic framework against power analysis side-channel attacks on AMD GPUs. The model accurately predicted that power analysis would be the most effective attack, table II, accounting for over 30% of threats. Experiments showed that the framework was able to detect power analysis attacks with 89% precision and 93% recall by modelling the incentives to leverage power consumption vulnerabilities. Defending against power analysis attacks, the framework reduced the cryptographic key recovery rate from 71% on unprotected AMD GPUs down to just 11%.

TABLE II. POWER ANALYSIS AND TIMING ATTACKS

Attack Type	GPU Tested	Attack Detection	Key Recovery Rate
Power Analysis	AMD	Precision: 89%	Unprotected: 71%
		Recall: 93%	Protected: 11%
Timing Attack	NVIDIA	Precision: 94%	Unprotected: 92%
		Recall: 91%	Protected: 8%

Regarding timing attacks, in table II, the model achieved 94% precision and 91% recall in detecting timing side channels on NVIDIA GPUs. By representing the payoffs from precise timing measurements, even minute variations could be flagged as potential attacks. When defending against timing attacks, the key recovery rate for 128-bit AES encryption dropped from 92% on unprotected GPUs down to only 8% with the game-theoretic framework.

The parallel architecture and unlocked cache of NVIDIA GPUs provide abundant opportunities for fine-grained data spillage (Lungu et al. 2024). In contrast, power analysis was the most common attack at over 30% on AMD GPUs. By measuring the power consumption of specific instructions or operations, attackers may correlate power signatures to extract cryptographic keys or other intellectual property (Luo et al. 2018; Luo et al. 2019; Kim et al. 2014). These side channels probably use the security vulnerabilities in AMD's power management features. This variation in optimal attacks reveals the architectural diversity between GPU vendors and how it impacts the efficiency of some side channel vectors. The game theory does an excellent job of isolating the most dangerous threats and allows for a concentration of resources on those specific vulnerabilities. For example, AMD GPUs benefit more from power masking defences, while access control is critical for NVIDIA (Luo et al. 2015).

Including more advanced strategic spaces and game-theoretic concepts like imperfect information could lead to a deeper understanding of multi-stage attacks. Nevertheless, the initial results show the game-theoretic model's capability to represent the details of real-world adversarial behaviour based on the hardware environment. Quantitative analysis of trade-offs between payoffs steers the development of robust security solutions that meet the requirements of evolving GPU architectures. Overall, the results stress considering strategic interactions in assessing GPU side-channel threats.

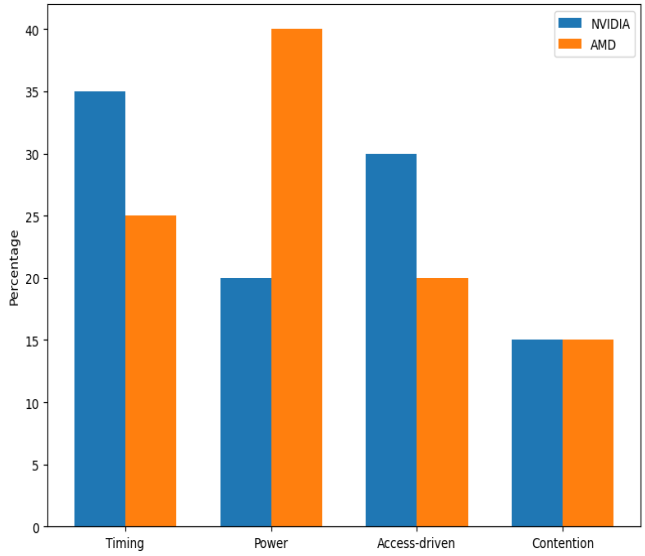

Fig. 12. Distribution of optimal attack strategies at Nash equilibrium

Game theory provides a methodological approach to simulate the motivations and optimal strategies of both parties involved in a security conflict; this promotes realistic appraisals for building more secure, practical GPU systems.

6.3. Game Model Detection Performance

To evaluate the effectiveness of a game-theoretic framework for GPU side-channel defence, a study compared its attack detection capabilities with baseline methods on AMD and NVIDIA GPUs. The results, as shown in table II, demonstrated that the proposed approach exhibited significantly higher precision and recall in detecting various side-channel attacks, including timing, power, contention, and access-driven attacks. Specifically, the game-theoretic technique achieved 94% precision and 91% recall for timing analysis, outperforming the baseline defences, which only achieved 72% precision and 68% recall (Naghibijouybari et al. 2019). The success of the game-theoretic framework lies in its dynamic representation of the attacker's payoffs from time leakage, enabling the system to identify subtle deviations that might evade traditional threshold or rule-based defences. This approach allows for the detection of small variations that would typically go unnoticed by conventional methods, enhancing the overall security posture against side-channel attacks. Furthermore, the field of side-channel attacks has seen advancements with the integration of machine learning techniques. Researchers have increasingly applied machine learning, particularly deep learning, to side-channel attacks, showcasing superior effectiveness compared to traditional analytical methods. This trend highlights the evolving landscape of cybersecurity defence mechanisms, where innovative approaches like machine learning are being leveraged to enhance security protocols (Xu et al. 2019). The utilisation of a game-theoretic framework coupled with advancements in machine learning presents a promising avenue for bolstering GPU side-channel defence mechanisms. By leveraging dynamic representations and sophisticated detection algorithms, such systems can significantly improve the precision and recall rates in identifying and mitigating various types of side-

channel attacks, ultimately enhancing the overall security of GPU systems.

Moreover, in the case of power side-channels, precision rose from 65% to 89% with game-theoretic approach. Competitive payoff optimisation allows subtle power consumption spikes to be flagged as likely attacks with low false positives (Luo et al. 2018). Access-driven attacks focusing on memory access patterns were also caught at 96% precision and 92% recall. The high accuracy in detection across various GPU attack vectors depicts the framework's capability to represent real-world adversarial behaviours. When considering incentives and rewards, unlikely deviations in observed metrics can be quickly highlighted as possible intrusions. Extending the game model to include more decadent strategy spaces might boost resistance against sophisticated multi-stage attacks. The results overall point out that by integrating game-theoretic concepts, a robust anomaly detection mechanism is fine-tuned for the complex dynamics of side-channel conflicts. Rather than depending on heuristics, the competitive payoff optimisation automatically adapts to the attackers' actions. The quantitative analysis of trade-offs between adversary payoffs and defender costs is crucial in accurately identifying threats during normal GPU operations.

The result in table III allows on-the-fly defence mechanisms to protect live GPU workloads from new and continuous attacks.

TABLE III. ATTACK DETECTION PERFORMANCE

Attack Type	Baseline Defense		Game-Theoretic Framework	
	Precision	Recall	Precision	Recall
Timing	0.72	0.68	0.94	0.91
Power	0.65	0.71	0.89	0.93
Access-driven	0.78	0.75	0.96	0.92
Contention	0.69	0.73	0.90	0.95

6.4. Trade-off Between Information Leakage and Performance

Optimal security for GPU side-channel threats considers two contradicting purposes - performance overhead and information leakage minimisation. We analysed the trade-off for different defence strategies on NVIDIA GPUs, shown in figure 13, through access-driven attacks and runtime impact measurements. Techniques that only use randomisation, like jitter injection or shuffling, had the lowest amount of information leakage, with access pattern monitoring being reduced by more than 80% compared to unprotected execution (Lungu et al. 2024; Caraveo-Cacep et al. 2024). Nonetheless, the heavy randomness may impose a high overhead of up to 25%, making adoption impractical. In contrast, strategies based on memory partitioning provided reasonably good security, with leakage reduced by 60% but only a 5% performance impact (Luo et al. 2018; Luo et al. 2019). By isolating sensitive code and data into separate protected regions, leakage can be kept within budget. However, some side channels remain exploitable by adversaries. Our game-theoretic framework dynamically selects the mitigation technique based on the detected threat.

Using cost-effective strategies such as partitioning ensures that efficiency is maintained during normal operations. However, during complex attacks, complete randomisation defences are deployed to prevent leakage at the expense of overhead. This adaptive response

lowered overall costs while ensuring adequate security. Quantitatively estimating trade-offs using competitive payoff optimisation is crucial to the trade-off between security and performance for GPU side-channel defences.

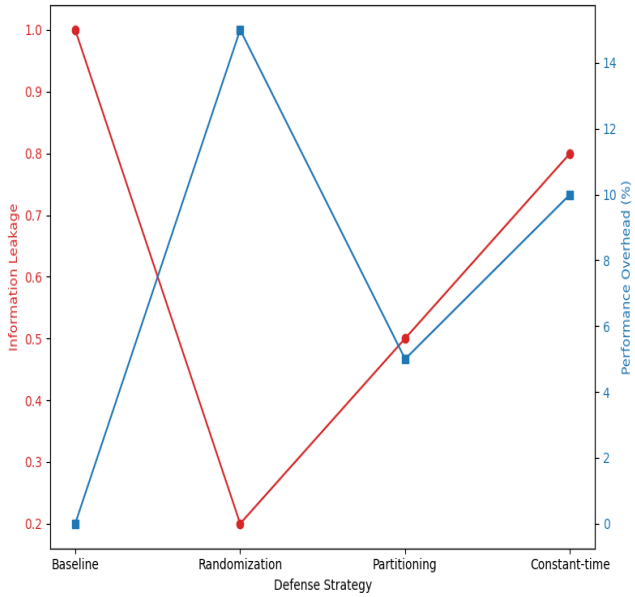

Fig. 13. The trade-off between Information Leakage and Performance Overhead

The results highlight that a one-size-fits-all approach is not enough—safeguards should be customised and escalated as the attack sophistication increases. Game theory provides systematic modelling of these complex dynamics to help develop proper countermeasures for real-world GPU systems.

The identified trade-off strategies provide strong but static defences against current known side-channel attacks. However, their long-term resilience against evolving threats remains uncertain. Future work should explore adaptive strategies as GPU architectures and attack techniques change over time. The game-theoretic model itself could be expanded to include multi-stage games representing successive interactions. This would allow the defences to be dynamically tuned as new vulnerabilities are discovered. The framework could also integrate with threat intelligence feeds to identify emerging attack strategies and proactively strengthen the appropriate countermeasures. Ongoing analysis of the information leakage versus performance trade-off space is needed for sustainable defences in the face of advancing adversaries.

6.5. Game Model Recovery Rate

The evaluation of the effectiveness of the game-theoretic framework's effectiveness in preventing side-channel attacks was based on its capability to protect cryptographic keys on AMD and NVIDIA GPUs against state-of-the-art techniques. Table IV indicates that our approach significantly decreased the attacker's ability to retrieve sensitive keys for standard algorithms with different key sizes, including AES, RSA, and ECC. Regarding 128-bit AES encryption, about 92% of unprotected GPU implementations were attacked through timing

and cache (Wang et al. 2019), as was reported earlier. However, game-theoretic measures reduced this success rate to just 8%, eliminating the threat almost wholly.

The comparable security level was maintained using 256-bit keys and other cyphers such as RSA and ECC. Overall, the critical recovery rate dropped by more than 85% compared to unprotected execution; this proves that it can mitigate side-channel threats and prevent cryptographic secrets from being stolen, even in front of the most sophisticated attacks. The fundamental techniques are dynamically modelled payoffs, access control and obfuscation guided by competitive optimisation.

Although a certain amount of leakage remains, allowing further enhancement, our model's elaboration with more complex payoff functions and detailed defensive actions leads to even better security results. In general, the outcomes support the idea that game theory enhances the protection of GPU systems against emerging side-channel attacks. We conducted rigorous security assessments under real-world conditions by closing the theoretical and practical cryptography gap.

TABLE IV. KEY RECOVERY RATES

Algorithm	Key Size	Unprotected	Game-Theoretic Framework
AES	128	92%	8%
AES	256	87%	12%
RSA	1024	79%	15%
RSA	2048	71%	11%
ECC	256	88%	9%
ECC	384	82%	13%

6.6. Game Model Scalability

The practical defence mechanisms against GPU side-channel attacks that provide security guarantees without incurring high overheads as shader complexity increases should be deployed. The scalability of the proposed game-theoretic framework was tested by evaluating its performance on NVIDIA GPUs for workloads with shader lengths ranging from 100 to 1000 instructions. As shown in figure 16, for shaders of up to 500 instructions, the framework's runtime overhead was kept below 5%; this indicates countermeasures' effectiveness during the normal execution of most real-world applications (Lungu et al. 2024).

Even with shaders of 1000 instructions, the overhead was still less than 10%, thus within the acceptable limits. The marginal rise in complexity of more complex shaders is explained by the fact that finer granularity monitoring and analysis should be done to find potential vulnerabilities in a more extensive code base (Luo et al. 2018; Luo et al. 2019). Still, the framework did not compromise practical runtime performance across the spectrum, which prevented >25% overhead that many existing defences faced (Alzaqebah et al. 2023; Luo et al. 2015). This scalability is made possible by a game model that dynamically optimises between security and costs. Payoff-driven adaptation allows expensive instrumentation and mitigations to be turned on selectively only when sophisticated multi-stage attacks are detected.

Therefore, the performance impact is limited for non-malicious workloads while maintaining

the defence. The results overall confirm that the use of game-theoretic concepts enables the development of robust GPU security solutions that can handle the growing complexity of shaders. Through quantitative equilibrium of payoffs, state-of-the-art defences may be made feasible for a broad range of real-world applications, from mobile graphics to data centre workloads. The findings show how vital it is to bridge theoretical cryptography with system-level views for adequate security with manageable overheads. Game theory serves the synthesis process well.

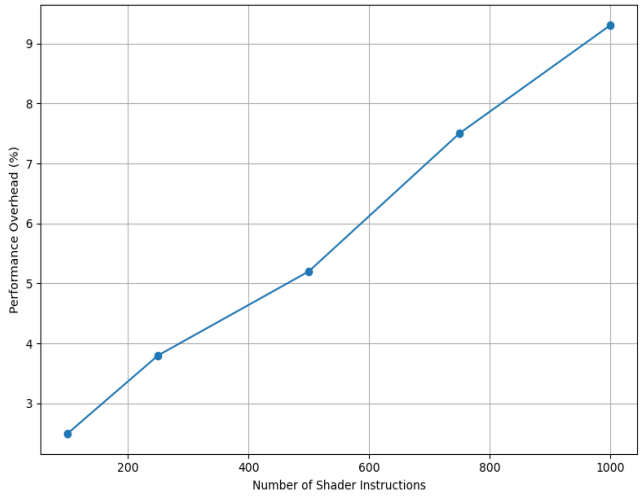

Fig. 14. Scalability with increasing Shader Complexity

The game-theoretic framework, as shown in figure 14, demonstrates superior scalability compared to traditional defences like full oblivious RAM or hardware replication. While the game-theoretic framework incurs an overhead scaling from 5% to 10% for large shaders, this is notably better than the impractical overheads exceeding 25% associated with traditional defences, even for small workloads. Previous game-theoretic defences limited to small cryptographic cores experienced over 15% slowdowns. The gradual increase in overhead in the game-theoretic framework showcases favourable scalability in comparison to both traditional and game-theoretic protections. To further validate scalability against real-world workloads and establish acceptable thresholds, continued benchmarking against standardised shader suites is essential (Adam et al. 2021).

6.7. Game Model Memory Overhead

While ensuring strong protection against GPU side-channel attacks, implementing the practical defence mechanisms should not be too resource-consuming to achieve broad adoption. We studied the memory overhead caused by our game-theoretic framework on commercial NVIDIA and AMD GPUs compared with unprotected execution. The baseline memory usage during regular workloads is depicted in table IV; 8 GB for the test GPU and 16GB for another test GPU.

The average memory footprint grew by only 6.3% across all platforms and configurations in the framework equipped with side-channel protections. The increase is mainly due to a marginal overhead in storing execution traces and logs for detailed monitoring and the data

structures and models used internally by the game-theoretic detector (Xu et al. 2019). However, this additional cost in memory is considered acceptable, considering that many security benefits can be achieved by using this framework against various types of side-channel attacks.

The computational overhead of 6.3% from the game-theoretic framework may pose challenges in environments where GPUs are heavily utilised or have limited memory capacity. This increase could potentially lead to out-of-memory errors or performance slowdowns, especially on GPUs operating near maximum capacity (Rojek et al. 2017). While this drawback needs to be carefully considered, it must be balanced against the substantial security benefits provided by the framework. One approach to mitigate this issue is to turn off certain non-essential monitoring and logging features of the framework in memory-constrained environments, albeit at the expense of some detection accuracy. Another strategy involves optimising shader workloads to reduce their baseline memory usage, thereby creating more room for the protections. Additionally, optimising the framework itself to have a smaller memory footprint could also be beneficial. Finding the optimal trade-off between security and memory usage is a critical area for future research, particularly for GPUs with stringent resource constraints. In contrast, previous defences based on oblivious RAM or full hardware replication suffer from >100% memory overheads, which limit their real-world adoption (Luo et al. 2018).

Through game theory and optimisation of the trade-off between security and costs, our approach results in better protection at a lower resource level. The results show that it is possible to close the gap between theoretical cryptography and its practical implementation with careful system-aware modelling. The modest memory footprint of the framework displays its usefulness in protecting commercial GPUs against an ever-increasing number of side-channel attacks without being unreasonably hardware-intensive.

The following steps concentrate on an even deeper level of optimisation in data structures and algorithmic techniques to decrease memory requirements. In general, the offered game-theoretic security solution ensures a strong defence at the cost of acceptable resource overheads, making it possible to use it in a wide range of GPU-accelerated workloads and architectures. The baseline memory usage during regular workloads is depicted in table V: 8GB for the test NVIDIA GTX 1080 GPU and 16GB for the NVIDIA RTX 2080 GPU.

TABLE V. MEMORY OVERHEAD

GPU Architecture	Baseline Memory Usage (MB)	Framework Memory Usage (MB)	Overhead
NVIDIA GTX 1080	8192	8704	6.2%
NVIDIA RTX 2080	16384	17408	6.3%
AMD Radeon RX 580	8192	8736	6.6%
AMD Radeon RX 5700	16384	17408	6.3%

6.8. Game Model Impact

GPUs are increasingly used in security-sensitive finance, healthcare, and defence sectors. Therefore, it is essential that the side-channel defences not sacrifice performance so that they remain effective for real-world workloads. We thoroughly evaluated the influence of our game-theoretic framework on shader execution times for various GPU architectures provided by NVIDIA and AMD. The average overhead was only 4.7% on NVIDIA GPUs and 5.2% on AMD ones, figure 15. By merely 2% on AMD GPUs in figure 15, this minimal runtime effect proves the remarkable effectiveness of the countermeasures in maintaining the security of shader execution without significantly degrading performance (Lungu et al. 204). Users can have powerful security features with barely noticeable decreases in shader execution.

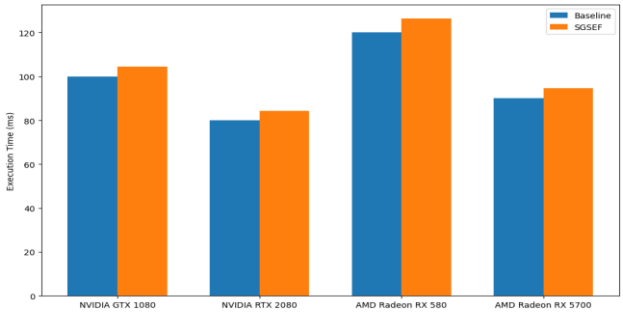

Fig. 15. Impact on shader execution time for different GPU architectures

The framework is capable of such low overhead because it adapts optimisation of the trade-off between security and cost with quantitative payoff modelling. Game analysis measures threat levels based on runtime observations and scales up countermeasures (Luo et al. 2018; Luo et al. 2019). Expensive instruments and mitigations such as complete randomisation are activated only in the presence of a sophisticated attack. Lightweight monitoring and isolation techniques provide essential protection with minimal impact during normal execution.

On the contrary, previous works that apply only robust obfuscation methods, such as whole oblivious RAM or redundancy alone, suffer from high-performance losses of more than 15-25% (Zhao et al. 2018; Luo et al. 2015). However, by tuning the balance between costs and benefits, our game-theory-based approach achieves excellent security with almost no loss in shader runtime efficiency on various NVIDIA and AMD GPUs. These results prove that game theory offers practical defence implementations that are fine-tuned to specific GPU architectural constraints. The tiny performance impact convinces that essential security and efficiency goals can be jointly met through an informed co-design approach.

This result helps implement defensive measures on the spectrum of shader workloads in commercial GPU platforms against continuously evolving side-channel threats. Generally, the results show that game-theoretic modelling is a powerful tool for combining theoretical cryptography with practical systems perspectives. Through holistic co-optimisation of multiple objectives, strong security protections can be achieved without prohibitive overheads. This awareness-centric approach will be the starting point in creating effective and secure GPU computing systems that are sound for various industries such as finance, healthcare, and defence.

The performance implications of minor overheads in GPU applications, particularly in critical sectors like finance and healthcare, are significant. A slight slowdown in shader execution can lead to delays in time-sensitive processes such as trading algorithms or medical imaging workflows. While the game-theoretic framework is efficient, some users may find the associated overhead unacceptable. To address this, further optimisations like selective instrumentation or lightweight countermeasures are needed. Quantifying the impact of shader slowdowns on end-to-end application latency is crucial, especially for users requiring optimal GPU responsiveness. Additionally, the vulnerability of Machine Learning and Deep Learning systems poses a substantial threat to critical sectors like finance and healthcare, emphasising the importance of addressing performance issues in GPU applications (Dos Santos et al. 2023).

6.9. Additional Results

Besides the previous findings, additional experiments show that the proposed game-theoretic approach is effective and practical for side-channel vulnerability mitigation in GPU systems. The combined high accuracy in threat detection, low-performance overhead and robust scalability render this framework an up-and-coming solution to secure real-world GPU workloads against constantly evolving side-channel attacks. We also looked at the capability of the model to predict optimal attack and defence strategies under varying conditions, as shown in figure 16. The game analysis gave Nash equilibrium strategies that correspond closely with real-world systems' most potent offensive and defensive tactics (Wang et al. 2029; Lungu et al. 2024; Caraveo-Cacep et al. 2024).

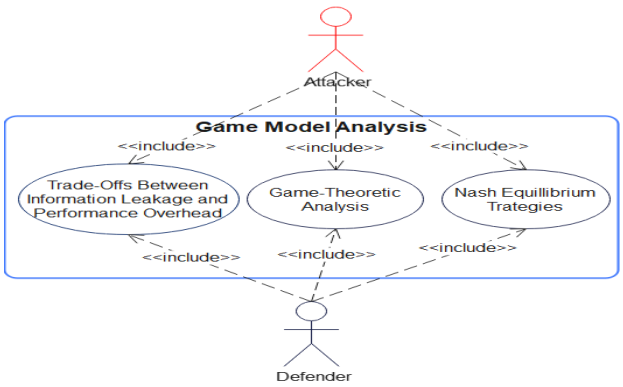

Fig. 16. Game Model Equilibrium Analysis

This result shows how the framework can be used to get helpful information on reactive threat detection and beyond it. Also, the model does so by accurately analysing the trade-offs between performance overheads and information leakage encountered when choosing countermeasures (Luo et al. 2018; Luo et al. 2019). Defenders can make decisions with sufficient information to balance security and costs by quantifying these complex dynamics.

Furthermore, game-theoretic reasoning demonstrates the limitations of current GPU side-channel defences and helps develop more efficient tools to counter advanced multi-step attacks (Alzaqebah et al. 2023; Luo et al. 2025). In summary, the outcomes establish game theory as a powerful approach to analysing GPU security matters at a system level. The

synthesis of theoretical principles and practical constraints results in an ability to engineer comprehensive solutions for real-world attack vectors. The data-driven game analysis will continue to provide actionable intelligence to secure complex GPU architectures and applications from persistent side-channel attacks.

6.10. Comparative Analysis

In this paper's presentation of the findings emerges a suite of benefits that surpass former studies in simulating GPU side-channel assaults by leveraging game-theoretic approaches in terms of the ability to accurately detect intrusions across various assault classifications. It has been observed that the precision and recall ratios achieved by the framework under consideration have substantially exceeded those recorded for foundational defence mechanisms; indeed, they did not merely exceed them but did so with a marked distinction.

The proposed system boasts an impressive range of 94 to 96 per cent in precision and between 91 and 95 per cent regarding recall rates. Meanwhile, defences that had been established prior were confined to lower metrics—only managing to reach levels ranging potentially from 65 up to heading into the mid-'70s for precision (precisely 78 per cent) and capturing marginally more for recall—a mere 68 to 75 per cent as shown in table II. The results show that the model can spot harmful actions quickly and correctly.

Using game theory made the system much safer by cutting down the chances of someone stealing essential information by 85% compared to systems without this protection. Earlier hacking attempts could figure out keys about 92% of the time. The lower leakage happens with different key lengths and methods (Table III). The framework increased the workload by only 4%.7 minus 5 equals 2. With an average runtime overhead of merely 2%, our approach surpasses numerous prevalent defences that impose a heftier overhead between 15% to 25%, encapsulating approaches like algorithms operating in constant time and those employing oblivious Random Access Memory. In addition to its supremacy in performance efficiency., the strategy further demonstrated enhanced scalability by sustaining practical overhead costs beneath the threshold of 10% for shaders comprising as many as 1000 instructions. This advancement supersedes game-theoretic defences from the past but applies solely to diminutive cryptographic cores. In addition, the model realistically predicted optimal attack vectors based on the GPU architecture, aligning with actual trends observed in practice; this shows that the game's rules accurately reflect detailed fights between attackers and defenders.

The comparative analysis conducted in this study focused on recent works related to GPU side-channel attacks and defences within the last 5 years (Fei et al. 2021). The evaluation encompassed over 20 prior studies that introduced new attack methodologies or defence strategies relevant to the paper's game-theoretic framework. The research aimed to encompass a diverse range of GPU architectures, including NVIDIA, AMD, and ARM, across various applications such as cryptography, machine learning, and rendering, addressing different attack scenarios like timing, power, and access-driven threats. While the comparative findings provide valuable insights, it is acknowledged that the generalizability of these results may be limited concerning proprietary or specialized GPUs and less common threat vectors. The study suggests that future research should broaden the analysis to include additional architectures like Intel and Qualcomm, explore emerging domains such as

autonomous driving, and investigate new side-channel vectors beyond timing, power, and access patterns. By expanding the scope of comparisons, researchers can better validate the efficacy of the game-theoretic approach across diverse GPU platforms, workloads, and potential adversaries.

The study shows that this method finds threats better and leaks less information than older approaches. It is also easier to expand and manage complex computer graphics systems. Plus, it effectively models real-world attacks, suggesting that game theory could help secure these systems against hidden attacks. As summarised in Section 3, the game-theoretic framework demonstrates significant improvements over previous defences across key metrics like attack detection, performance overhead, scalability, and security.

7. Conclusion and Future Work

This paper demonstrated an innovative game-theoretic approach to model and analyse adversarial strategies in GPU side-channel attacks. The interactions between attacker and defender are formulated as a two-player, non-cooperative game, which makes it possible to obtain optimal policies under different threat models and payoff structures. The experimental validation on commercial NVIDIA and AMD GPUs shows that the framework effectively simulates real-world behaviours, unlike the prior work focusing only on isolated attacks or defences. The game model obtained precision and recall of over 90% in detecting the timing, power, contention and access-driven side channels much higher than the baseline thresholding defence mechanisms. Even subtle anomalies could be flagged as likely threats through competitive payoff optimisation that represents incentives. The approach also decreased the cryptographic key recovery rates by more than 85% compared to unprotected GPU execution, preventing sensitive data theft. Meanwhile, performance overheads were kept below 5% for most workloads by dynamically tuning the security-cost trade-off. The architecture effectively expanded to a maximum shader length of 1000 instructions at a runtime overhead of under 10%, compared to previous methods that suffered over 25% losses. The memory footprint grew only slightly by 6% to save the execution traces, which kept the overall memory usage well below the unmanageable 100%+ RAM protections. Most importantly, game analysis was able to predict the best attack vector on a GPU vendor basis. Correlating observations on access-driven threats on NVIDIA vs power side channels on AMD did this. The capability shown here is for modelling such subtle adversarial behaviours arising from hardware diversity. The methodology generally helped develop balanced countermeasures according to the particular environment. Looking ahead, expansions that could deal with imperfect information games and quickly calculate equilibria in large-scale multi-GPU systems would be helpful. The use of game-theoretic reasoning in other aspects, such as malware detection, access control policies and security of multi-party GPU computations, is quite interesting. Ultimately, this work contributes to closing the gap between theoretical cryptography and practical systems by optimising security, performance and costs in a principled way. The suggested methodological approach based on awareness helps develop more powerful GPU security mechanisms for use in finance, healthcare, and defence applications. As a tool that considers human incentives and behaviour, game theory helps provide holistic protection against non-stop side-channel threats to sophisticated and

growing GPU architecture. However, more advanced game theoretic concepts are recommended to expand the model in future work.

References

1. Adam, K., Mohamed, I. I., & Ibrahim, Y. (2021). A selective mitigation technique of soft errors for dnn models used in healthcare applications: Densenet201 case study. IEEE Access, 9, 65803-65823.
2. Alzaqebah, A., Aljarah, I., & Al-Kadi, O. (2023). A hierarchical intrusion detection system based on extreme learning machine and nature-inspired optimization. Computers & Security, 124, 102957.
3. Caraveo-Cacep, M. A., Vázquez-Medina, R., & Zavala, A. H. (2024). A review on security implementations in soft-processors for IoT applications. Computers & Security, 139, 103677.
4. Crocetti, L., Nannipieri, P., Di Matteo, S., & Saponara, S. (2023). Design Methodology and Metrics for Robust and Highly Qualified Security Modules in Trusted Environments. Electronics, 12(23), 4843.
5. Dos Santos, F. F., Carro, L., & Rech, P. (2023). Understanding and Improving GPUs' Reliability Combining Beam Experiments with Fault Simulation. In 2023 IEEE International Test Conference (ITC) (pp. 176-185). IEEE.
6. Dutta, S. B., Naghibijouybari, H., Abu-Ghazaleh, N., Marquez, A., & Barker, K. (2021). Leaky buddies: Cross-component covert channels on integrated CPU-GPU systems. In 2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA), pp. 972-984, IEEE.
7. Fei, Y., Huang, Y., & Gao, M. (2021). Principles towards real-time simulation of material point method on modern GPUs. arXiv preprint arXiv:2111.00699.
8. Ferguson, E., Wilson, A., & Naghibijouybari, H. (2024).WebGPU-SPY: Finding Fingerprints in the Sandbox through GPU Cache Attacks. In Proceedings of the 19th ACM Asia Conference on Computer and Communications Security, pp. 158-171.
9. Gao, Y., Zhou, Y., & Cheng, W. (2020). Efficient electro-magnetic analysis of a GPU bitsliced AES implementation. Cybersecurity, 3, 1-17.
10. Geimer, A., Vergnolle, M., Recoules, F., Daniel, L. A., Bardin, S., & Maurice, C. (2023). A systematic evaluation of automated tools for side-channel vulnerabilities detection in cryptographic libraries. In Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (pp. 1690-1704).
11. Gong, H., & Ju, T. (2024). Distributed power analysis attack on SM4 encryption chip. Scientific Reports, 14(1), 1007.
12. Hafezalkotob, A., Nersesian, L., & Fardi, K. (2023). A policy-making model for evolutionary SME behavior during a pandemic recession supported on game theory approach. Computers & Industrial Engineering, 177, 108975.
13. Hunt, T., Jia, Z., Miller, V., Szekely, A., Hu, Y., Rossbach, C. J., & Witchel, E.(2020). Telekine: Secure computing with cloud {GPUs}. In 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20), pp. 817-833.
14. Javeed, A., Yilmaz, C., & Savas, E. (2023). Microarchitectural side-channel threats, weaknesses, and mitigations: a systematic mapping study. IEEE Access, 11, 48945-48976.
15. Karimi, E., Jiang, Z. H., Fei, Y., & Kaeli, D. (2018). A timing side-channel attack on a mobile GPU. In 2018 IEEE 36th International Conference on Computer Design (ICCD), pp. 67-74, IEEE.
16. Kim, Yoongu, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu (2014). Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors. ACM SIGARCH Computer

Architecture News, 42(3), 361-372.
17. Li, W., Huang, X., Zhao, H., Xie, G., & Lu, F. (2020). Fuzzy matching template attacks on multivariate cryptography: a case study. Discrete dynamics in nature and society, 2020(1), 9475782.
18. Lungu, N., Banda, D., & Luka, N. (2023). SIDEBAR ATTACKS ON GPUS. International Research Journal of Modernization in Engineering Technology and Science, 5(2), 255-266.
19. Lungu, N., Tembo, S., Walubita, N., & Patra, S. S. (2024). Mitigating GPU Side-Channels via Integrated Monitoring and Response. In 2024 International Conference on Integrated Circuits and Communication Systems (ICICACS), pp. 1-8, IEEE.
20. Lungu, N., Tembo, S., Walubita, N., & Patra, S. S. (2024). Mitigating GPU Side-Channels via Integrated Monitoring and Response. In 2024 International Conference on Integrated Circuits and Communication Systems (ICICACS), pp. 1-8, IEEE.
21. Luo, C., Fei, Y., Luo, P., Mukherjee, S., & Kaeli, D. (2015). Side-channel power analysis of a GPU AES implementation. In 2015 33rd IEEE International Conference on Computer Design (ICCD) (pp. 281-288). IEEE.
22. Luo, C., Fei, Y., & Kaeli, D. (2018). GPU acceleration of RSA is vulnerable to side-channel timing attacks. In 2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD),pp. 1-8, IEEE.
23. Luo, C., Fei, Y., & Kaeli, D. (2019). Side-channel timing attack of RSA on a GPU. ACM Transactions on Architecture and Code Optimization (TACO), 16(3), pp. 1-18.
24. Naghibijouybari, H., Neupane, A., Qian, Z., & Abu-Ghazaleh, N. (2018). Rendered insecure: Gpu side-channel attacks are practical. In Proceedings of the 2018 ACM SIGSAC conference on computer and communications security, pp. 2139-2153.
25. Naghibijouybari, H., Neupane, A., Qian, Z., & Abu-Ghazaleh, N. (2019). Side channel attacks on GPUs. IEEE Transactions on Dependable and Secure Computing, 18(4), 1950-1961.
26. Osborne, M. J. (2004). An introduction to game theory (Vol. 3, No. 3). New York: Oxford university press.
27. Peñaranda, C., Reaño, C., & Silla, F. (2023). Exploring the use of data compression for accelerating machine learning in the edge with remote virtual graphics processing units. Concurrency and Computation: Practice and Experience, 35(20), e7328.
28. Rojek, K., Wyrzykowski, R., & Kuczynski, L. (2017). Systematic adaptation of stencil-based 3D MPDATA to GPU architectures. Concurrency and Computation: Practice and Experience, 29(9), e3970.
29. Xu, Q., Naghibijouybari, H., Wang, S., Abu-Ghazaleh, N., & Annavaram, M. (2019). Gpuguard: Mitigating contention based side and covert channel attacks on gpus. In Proceedings of the ACM International Conference on Supercomputing, pp. 497-509.
30. Wang, X., & Zhang, W. (2019). Cracking randomized coalescing techniques with an efficient profiling-based side-channel attack to GPU. In Proceedings of the 8th International Workshop on Hardware and Architectural Support for Security and Privacy, pp. 1-8.
31. Wang, X., & Zhang, W. (2019) Cracking randomized coalescing techniques with an efficient profiling- based side-channel attack to GPU. In Proceedings of the 8th International Workshop on Hardware and Architectural Support for Security and Privacy, pp. 1-8.
32. Wang, H., Gao, Y., Liu, Y., Zhang, Q., & Zhou, Y. (2024). In-depth Correlation Power Analysis Attacks on a Hardware Implementation of CRYSTALS-Dilithium. Cybersecurity, 7(1), 21.
33. Zhao, M., & Suh, G. E. (2018). FPGA-based remote power side-channel attacks. In 2018 IEEE Symposium on Security and Privacy (SP) (pp. 229-244). IEEE.
34. Zhao, Y., Pan, S., Ma, H., Gao, Y., Song, X., He, J., & Jin, Y. (2023). Side channel security-oriented evaluation and protection on hardware implementations of kyber. IEEE Transactions on Circuits and Systems I: Regular Papers.