Global Multidisciplinary Journal
eISSN: 2791-173X pISSN: 2791-2760

VOLUME04 ISSUE09
Published 10-09-2025 *Page No. 42-50*

Towards Secure, Trusted, and Virtualized Multi-Tenant FPGA–Cloud Ecosystems: A Comprehensive Research Framework Integrating Hardware Roots of Trust, Cryptographic Acceleration, and Zero-Trust Cloud Security

Arvind Raman

Department of Computer Science, University of Wellington, New Zealand

ABSTRACT

This research article provides a comprehensive and integrative examination of security, trust, virtualization, and cryptographic enablement in multi-tenant cloud environments incorporating Field-Programmable Gate Arrays (FPGAs). Drawing on a diverse range of foundational and contemporary studies, the article synthesizes architectural, cryptographic, and policy-driven security concerns across hardware-based trust mechanisms, FPGA virtualization, secure data retrieval, cloud adoption, and zero-trust methodologies. The growing reliance on heterogeneous compute infrastructures, particularly the integration of reconfigurable hardware into cloud platforms, has intensified security challenges such as remote physical attacks, multi-tenant inference leakage, bitstream manipulation, data privacy risks, and trust management failures. The review unifies disparate areas including TrustZone-based system isolation, Trusted Platform Module (TPM) authorization, Physically Unclonable Function (PUF) protections, homomorphic encryption acceleration, multi-tenant risk vectors in reconfigurable hardware, and SaaS-level tenant isolation. It further integrates advanced cloud data-security frameworks encompassing secure attribute-based sharing, hybrid elliptic-curve cryptography, hierarchical indices for encrypted search, and secure auditing via Diffie–Hellman-based schemes. Methodologically, the research develops a conceptual synthesis that reinterprets existing findings through a multi-layered security lens grounded in zero-trust architectural principles. The results highlight persistent vulnerabilities in multi-tenant FPGA clouds, persistent governance shortcomings, inconsistencies in cryptographic enforcement models, and limitations in current virtualization stacks. It also identifies emergent opportunities, particularly in homomorphic-encryption-enabled federated learning, energy-aware distributed cloud security, and trusted FPGA provisioning for heterogeneous environments. The discussion proposes an expanded zero-trust FPGA-cloud model emphasizing continuous attestation, cryptographic binding of hardware identities, granular tenant isolation, and resilience against side-channel and remote physical manipulations. The article concludes that future secure FPGA–cloud ecosystems must be architected around hardware-anchored trust, dynamic policy-driven cryptography, and full-stack multi-tenant isolation integrated into cloud orchestration frameworks.

KEYWORDS

FPGA cloud security; multi-tenancy; hardware trust; virtualization; zero-trust security; encrypted cloud data; homomorphic encryption

INTRODUCTION

The evolution of cloud computing from monolithic virtual machines toward heterogeneous, distributed, and reconfigurable infrastructures presents a sophisticated array of security challenges that surpass traditional software-only protection paradigms. Early studies conceptualized cloud computing primarily through virtualized resource pools abstracted from physical hardware (Marinescu, 2013). As cloud ecosystems

expanded, the adoption frameworks devised for enterprise and business clouds emphasized security as the determinant of sustainable cloud adoption (Chang, Kuo & Ramachandran, 2016). However, this shift toward large-scale, multi-tenant computational fabrics soon introduced unprecedented concerns related to isolation, confidentiality, and trust management (Gartner, 2011). These concerns intensified with the pervasive integration of the Internet of Things (IoT), which aggregated massive distributed data pipelines into cloud back-ends and magnified the attack surface through device heterogeneity and cross-layer dependencies (Botta et al., 2016).

Recent cloud trends, including distributed cloudlets, fog computing, and energy-aware offloading models, introduced additional dimensions where security remains deeply intertwined with performance, energy efficiency, and resource allocation (Gai et al., 2016; Jalali et al., 2016; Varghese & Buyya, 2018). As these models matured, new paradigms involving encrypted data retrieval, secure data sharing, privacy-preserving analytics, and multi-keyword encrypted search began reshaping cloud architectures to emphasize cryptographic assurances and tenant-centric security boundaries (Ali et al., 2020; Kumar & Bhatt, 2020; Swami & Das, 2020; Vengala, Kavitha & Kumar, 2020; Li et al., 2018). These developments collectively underscore a long-standing gap: securing multi-tenant cloud systems requires protections grounded not solely in software isolation but increasingly in hardware-anchored security primitives.

One of the most significant hardware-level shifts has been the widespread integration of FPGA accelerators in cloud platforms. Driven by demands for low-latency computing, customizable logic, and performance acceleration, cloud providers have deployed FPGAs to support diverse applications including machine learning, cryptography, scientific simulations, and high-performance networking. However, the reconfigurable nature of FPGAs introduces new vulnerabilities absent in conventional CPU-only environments (Zhang & Qu, 2019). Multi-tenant FPGA infrastructures are particularly exposed to bitstream manipulation, hardware Trojan insertion, timing-based side-channel attacks, and resource-contention attacks that exploit shared logic fabric states (Zeitouni, Dessouky & Sadeghi, 2020).

Virtualization of FPGAs, once envisioned to provide strong isolation and flexible provisioning, has itself become a complex source of security risk. Architectures like virtualized logic clusters, overlay architectures, and hypervisor-controlled reconfiguration interfaces have increased system efficiency but simultaneously expanded attack vectors. Research on FPGA virtualization and cluster-scale orchestration, such as the work on Hetero-ViTAL and cloud-scale FPGA resource abstraction, illustrates both the opportunities and threats associated with sharing reconfigurable hardware among untrusted tenants (Zha & Li, 2020; Zha & Li, 2021). Complementary work on trusted hardware mechanisms such as TrustZone integration within Zynq SoCs and TPM-based authorization demonstrates the emerging convergence between hardware-rooted trust and secure virtualization (Gosain & Palanichamy, 2014; Yu et al., 2016).

Alongside these developments, attack-specific research has revealed vulnerabilities such as remote physical attacks on FPGAs that exploit timing-difference circuits, with defenses such as DARPT designed to mitigate these threats (Zhang et al., 2022). Furthermore, cryptographic innovations such as homomorphic-encryption acceleration implemented directly on FPGAs offer promising solutions for federated learning and privacy-preserving cloud computation, yet these systems demand high assurance in the underlying hardware platform to prevent adversarial misuse or inference-leakage (Yang, Hu & Chen, 2020).

The culmination of these developments has led to the emergence of zero-trust models for cloud environments, particularly emphasizing continuous authentication, strict least-privilege access, real-time monitoring, and the assumption that no component—whether user, device, or hardware accelerator—should be inherently trusted

(Hariharan, 2025). This shift aligns with the contemporary realization that securing modern cloud ecosystems requires multi-layered, adaptive, and hardware-aware protection mechanisms capable of confronting the full spectrum of threats from malicious tenants to compromised firmware.

Despite substantial progress, a clear literature gap persists: a comprehensive, hardware-driven security architecture for multi-tenant FPGA clouds, tightly integrated with cryptographic access controls and zero-trust frameworks, has not yet materialized. Existing approaches either focus narrowly on hardware attack mitigation, or on cloud-level data security, but rarely address the architectural synthesis required for secure, large-scale FPGA integration in multi-tenant environments.

This article fills that gap by synthesizing research across hardware trust mechanisms, FPGA virtualization, multi-tenant risk models, encrypted cloud computing, and zero-trust frameworks. It constructs a unified conceptual model for secure FPGA-cloud ecosystems and presents a detailed analysis of methodological, architectural, and security implications. Through extensive theoretical elaboration, the study positions hardware-rooted trust and dynamic cryptographic controls as the foundation for next-generation secure cloud infrastructures.

METHODOLOGY

The methodological framework of this research is grounded in a multi-layered synthesis approach that draws on architectural, cryptographic, and policy-level studies from diverse research domains. Instead of employing empirical measurements, formal proofs, or simulation-based evaluations, the methodology relies on integrative theoretical reasoning designed to unify disparate threads of cloud security, FPGA virtualization, and hardware trust. This approach is justified by the scope of the research problem: multi-tenant FPGA cloud security is not confined to a single technical layer but emerges from interactions across hardware, software, virtualization stacks, and security governance. The method aims to reconstruct a conceptual model of these interactions by examining the implicit relationships, assumptions, and dependencies revealed in existing research.

The first methodological step involves categorizing the references into thematic clusters aligned with broader security architecture domains. These clusters include:

1. Hardware trust and platform security, including TrustZone technologies, TPM security analysis, and PUF-based IP protection (Gosain & Palanichamy, 2014; Yu et al., 2016; Jiliang Zhang et al., 2015).
2. FPGA attack surfaces and mitigation, covering multi-tenant threats, remote physical attacks, and general FPGA system vulnerabilities (Zeitouni et al., 2020; Zhang et al., 2022; Zhang & Qu, 2019).
3. Virtualization frameworks for reconfigurable hardware, including cloud FPGA virtualization and cluster-scale orchestration models (Zha & Li, 2020; Zha & Li, 2021).
4. Cryptographic data-security frameworks in the cloud, including attribute-based encryption, ECC-based hybrid schemes, searchable encryption, homomorphic encryption acceleration, and secure access control (Kumar & Bhatt, 2020; Li et al., 2018; Indhuja et al., 2017; Yang et al., 2020).
5. Cloud computing governance, adoption, and multi-tenant policy frameworks, including cloud security guidance, adoption models, and virtualization risk evaluations (Chang et al., 2016; CSA, 2011; Gartner, 2011).
6. Zero-trust security models, emphasizing continuous authentication and the elimination of implicit trust (Hariharan, 2025).

By synthesizing across these clusters, the methodology constructs an interpretive model that explores how hardware-based trust mechanisms interact with cloud-level cryptographic controls and zero-trust governance

frameworks.

The second methodological step involves a detailed conceptual mapping process. This mapping reconstructs the dependencies between cloud-level encryption requirements, FPGA virtualization barriers, secure attestation systems, and physical attack mitigation. The method prioritizes textual cross-examination, highlighting theoretical tensions such as the contrast between reconfigurability and isolation, performance optimization versus strong encryption, tenant convenience versus hardware-enforced restrictions, and the inherent contradictions between legacy trust models and contemporary zero-trust assumptions.

The third methodological step consists of theoretical extrapolation, whereby existing findings are extended to propose new synthesis principles. For example, integrating PUF-based identity binding with zero-trust attestation frameworks, or connecting homomorphic-encryption FPGA acceleration to secure federated learning governance. The goal is not to claim empirical verification but to demonstrate conceptual coherence between hardware trust and cloud cryptographic enforcement mechanisms.

The fourth methodological step involves the construction of a layered security architecture model that aligns FPGA hardware protections with cloud-level controls. This involves extensive argumentation explaining how hardware boundaries, virtualization containers, encryption models, and governance policies can reinforce each other under a unified zero-trust FPGA-cloud framework.

Finally, the methodology employs critical analysis to evaluate the limitations of current approaches and derive future research opportunities. The emphasis throughout is on comprehensive elaboration, conceptual detail, and theoretical rigor, ensuring that the proposed model advances beyond existing fragmented perspectives.

RESULTS

The synthesis reveals several critical insights into the nature of security in multi-tenant FPGA-cloud ecosystems. These findings emerge from a deep comparative analysis of hardware trust mechanisms, FPGA reconfigurability, cloud cryptographic frameworks, and zero-trust doctrines.

First, the examination of hardware trust primitives demonstrates that secure multi-tenant FPGA clouds must rely on hardware-anchored security roots rather than purely software-defined protections. Technologies like TrustZone within Zynq SoCs enable isolated execution environments by splitting the system into secure and non-secure domains, providing a foundation for trusted boot, secure monitor calls, and controlled access to resources (Gosain & Palanichamy, 2014). TPM-based authorization mechanisms add an additional layer by securing access control operations through cryptographic integrity checks and HMAC-based authorization (Yu et al., 2016). Moreover, PUF-based identity binding creates a hardware-unique signature that can support secure licensing and IP protection (Jiliang Zhang et al., 2015). The integrated outcome reveals that hardware-rooted trust is essential for managing FPGA tenants because reconfigurable logic cannot be reliably protected using software isolation alone.

Second, the analysis uncovers pervasive vulnerabilities in multi-tenant FPGA environments. Side-channel attacks, bitstream manipulation, timing-based information leakage, and resource contention attacks pose substantial threats (Zhang & Qu, 2019). Multi-tenant FPGA studies reveal that adversaries can observe power draw, timing behaviour, or residual configuration states to infer other tenants' computations (Zeitouni et al., 2020). Even more concerning are remote physical attacks that exploit hardware timing-difference circuits, as demonstrated in the DARPT study, illustrating that remote attackers can induce or exploit physical effects without physical access, significantly expanding the attack surface (Zhang et al., 2022). This evidence shows that virtualization alone cannot compensate for FPGA-specific risks.

Third, the research reveals that FPGA virtualization frameworks introduce both opportunities and new risks. Cloud-scale virtualization stacks such as those described in the ViTAL and Hetero-ViTAL architectures offer dynamic allocation, workload migration, and tenant isolation through partial reconfiguration and virtual logic overlays (Zha & Li, 2020; Zha & Li, 2021). However, these virtualization layers struggle to fully separate tenants at the hardware level because shared interconnects, routing structures, and power distribution networks remain exposed. Virtualization also demands highly trusted orchestration software because the hypervisor becomes a privileged intermediary controlling bitstream delivery and FPGA access, echoing concerns raised in earlier research such as CloudVisor, where nested virtualization was introduced to protect VMs from compromised hypervisors (Zhang, Chen & Zang, 2011). Translating this to FPGA clouds suggests that FPGA hypervisors themselves represent a critical security bottleneck.

Fourth, the evaluation of encrypted cloud computing frameworks demonstrates the necessity of coupling hardware-based security with sophisticated cryptographic methods. Hybrid ECC-based encryption enhances confidentiality for multi-tenant data flows (Kumar & Bhatt, 2020). Attribute-based encryption enables granular access control for resource-limited cloud users (Li et al., 2018). Diffie–Hellman-based auditing provides efficient and secure verification of cloud storage integrity (Yarava & Singh, 2019). Hierarchical clustering indices facilitate secure multi-keyword search on encrypted datasets (Indhuja et al., 2017). These findings collectively indicate that tenants rely increasingly on advanced cryptography to maintain data privacy. However, deploying these cryptographic tools in FPGA-accelerated workflows requires trustworthy hardware, emphasizing the interdependence between FPGA trust and cloud data privacy.

Fifth, homomorphic-encryption acceleration on FPGAs presents an emerging opportunity but also reveals the necessity of robust trust boundaries. Homomorphic encryption enables federated learning without exposing raw data to cloud servers, and FPGA acceleration reduces its computational cost (Yang et al., 2020). Yet the privacy benefits of homomorphic encryption degrade if the FPGA infrastructure itself is compromised. Therefore, secure homomorphic-encryption acceleration depends on both cryptographic correctness and hardware trust.

Sixth, cloud governance frameworks highlight persistent gaps in cloud security policy. Cloud Security Alliance guidelines emphasize multi-layered protection but often overlook hardware-specific risks (CSA, 2011). Cloud adoption frameworks emphasize business-level concerns but leave tenant-specific FPGA protections undefined (Chang et al., 2016). Studies on virtualization security emphasize the risks of shared hardware but do not fully account for reconfigurable-hardware-specific weaknesses (Gartner, 2011). The results highlight a misalignment between cloud governance and FPGA-infrastructure security.

Seventh, the integration of zero-trust principles into FPGA-cloud governance emerges as a compelling and necessary direction. Zero-trust eliminates implicit trust in tenants, users, hardware, and workloads, mandating continuous authentication and verification (Hariharan, 2025). Mapping these principles onto FPGA clouds results in a model where hardware trust, dynamic cryptographic controls, and continuous monitoring converge to mitigate multi-tenant FPGA threats.

DISCUSSION

The theoretical synthesis developed in this research underscores the inadequacy of isolated or single-layer protections in securing multi-tenant FPGA cloud environments. Instead, a comprehensive security architecture must integrate hardware-rooted trust, cryptographic access control, secure virtualization, and zero-trust governance. This discussion interprets the findings through a deeper conceptual lens, exposing tensions,

limitations, and emerging opportunities for future research.

A major tension arises between the reconfigurability of FPGAs and the predictability demanded by strong security boundaries. Reconfigurability inherently allows users to load custom logic, modify hardware-level behaviour, and generate new routing configurations on demand. While this flexibility enhances performance and adaptability, it simultaneously undermines deterministic control. For example, shared routing resources and global clock networks cannot easily be partitioned without reducing the hardware's programmable flexibility (Zeitouni et al., 2020). Even virtual overlays, such as those introduced in FPGA virtualization frameworks, cannot fully conceal underlying physical structures from sophisticated attackers. This tension suggests that multi-tenant sub-FPGA partitioning contradicts essential FPGA design principles, requiring future research to re-engineer FPGA fabrics with security-aware partitioning.

Another major tension exists between high-performance cryptographic operations and the computational overhead required for secure key management, access control, and encryption enforcement. Hybrid encryption systems (Kumar & Bhatt, 2020), attribute-based encryption (Li et al., 2018), and secure multi-keyword search frameworks (Indhuja et al., 2017) add substantial processing overhead. When executed on FPGAs, these operations compete with application-level logic for resource utilization, heightening the risk of information leakage through power, timing, or resource contention channels (Zhang & Qu, 2019). Therefore, cryptographic frameworks must be redesigned to avoid exacerbating hardware side-channel risks.

The discussion also highlights limitations in current trust mechanisms. For example, TrustZone provides secure worlds but does not protect the FPGA fabric itself when used in heterogeneous SoCs (Gosain & Palanichamy, 2014). TPM-based authorization secures access to sensitive operations but cannot prevent FPGA-specific side-channel attacks (Yu et al., 2016). PUF-based protection ensures identity integrity but does not mitigate FPGA resource-contention risks (Jiliang Zhang et al., 2015). The combined implication is clear: existing trust technologies serve as foundational components but fail to address the full spectrum of FPGA-cloud vulnerabilities.

A further limitation lies in the absence of cohesive policy frameworks governing FPGA virtualization and tenant isolation. While cloud providers implement access logs, cryptographic channels, and VM-level isolation, they seldom apply comparable controls at the FPGA fabric level. The lack of standardized policies for FPGA bitstream validation, tenant grouping, workload profiling, or suspicious-pattern detection results in fragmented risk management. The Cloud Security Alliance guidelines, while extensive, fail to address reconfigurable hardware threats (CSA, 2011). This indicates a significant governance gap.

Zero-trust principles offer a framework to reconcile many of these tensions. The zero-trust doctrine assumes that no user, device, or hardware module is inherently trustworthy (Hariharan, 2025). Applying this to multi-tenant FPGAs leads to several implications:

1. FPGA bitstreams must undergo continuous verification rather than one-time validation.

2. Tenant workloads must be monitored for anomalous electrical or timing behaviour that may indicate malicious activity.

3. Access to FPGA programming interfaces must be authenticated at every operation, not merely at session start.

4. Hardware-level attestations must be cryptographically bound to PUF signatures or TPM measurements.

5. FPGA virtualization hypervisors must enforce dynamic isolation rather than static partitioning.

The theoretical contribution of this research lies in recognizing that integrating zero-trust principles with

hardware-rooted trust and cryptographic data controls results in a robust full-stack security architecture for FPGA-cloud systems.

Future research opportunities arising from this discussion include:

- designing FPGA fabrics that inherently support secure multi-tenant partitioning;
- embedding real-time hardware monitoring systems that detect side-channel anomalies;
- integrating homomorphic-encryption acceleration with secure attestation systems;
- developing cloud orchestration systems that dynamically adjust isolation boundaries based on behavioural analytics;
- creating standardized governance frameworks specifically for cloud-hosted reconfigurable hardware.

The synthesis reveals that the future of FPGA-cloud security lies at the intersection of hardware innovation, cryptographic modernization, and zero-trust re-architecting of cloud governance.

CONCLUSION

This research provides a comprehensive consolidation of multi-tenant FPGA-cloud security by synthesizing key findings across hardware trust mechanisms, virtualization frameworks, cloud-level cryptographic protections, and zero-trust governance models. The results demonstrate that secure FPGA-cloud ecosystems require an integrated approach combining hardware-anchored trust, dynamic cryptographic protections, enforced tenant isolation, and continuous verification mechanisms. FPGA-specific vulnerabilities—ranging from physical side-channel attacks to tenant cross-inference—cannot be mitigated through traditional cloud security frameworks alone. Instead, zero-trust principles offer a strategic path forward by eliminating implicit trust within hardware accelerators, bitstream flows, and cloud orchestration layers. The study identifies multiple future research directions focusing on secure FPGA architecture redesign, cryptographic-hardware integration, anomaly detection, and unified cloud governance frameworks. The overarching conclusion is that multi-tenant FPGA-cloud systems will remain insecure without strong hardware trust foundations and dynamic, policy-driven, zero-trust security enforcement.

REFERENCES

1. Ali, F. S., Saad, H. N., Sarhan, F. H., and Naaeem, B. Enhance manet usability for encrypted data retrieval from cloud computing. Indonesian Journal of Electrical Engineering and Computer Science, 18, 2020.
2. Botta, A., De Donato, W., Persico, V., and Pescapé, A. Integration of cloud computing and internet of things: a survey. Future Generation Computer Systems, 56, 2016.
3. Cai, H., Wang, N., and Zhou, M. J. A transparent approach of enabling SaaS multi-tenancy in the cloud. IEEE 6th World Congress on Services, 2010.
4. Chang, V., Kuo, Y.-H., and Ramachandran, M. Cloud computing adoption framework: A security framework for business clouds. Future Generation Computer Systems, 57, 2016.
5. Cloud Security Alliance (CSA). Security Guidance for Critical Areas of Focus in Cloud Computing V3.0, 2011.
6. Gai, K., Qiu, M., Zhao, H., Tao, L., and Zong, Z. Dynamic energy-aware cloudlet-based mobile cloud computing model for green computing. Journal of Network and Computer Applications, 59, 2016.

7. Gartner Inc. Six Most Common Virtualization Security Risks and How to Combat Them. 2011.

8. Gosain, Y., and Palanichamy, P. TrustZone technology support in Zynq-7000 all programmable SoCs. Xilinx White Paper, 2014.

9. Hariharan, R. Zero trust security in multi-tenant cloud environments. Journal of Information Systems Engineering and Management, 2025.

10. Indhuja, A., Shaik, R. B. M. V., and Sujatha, P. A multi-keyword ranked search scheme over encrypted data based on hierarchical clustering index. International Journal on Smart Sensing and Intelligent Systems, 10, 2017.

11. Jalali, F., Hinton, K., Ayre, R., Alpcan, T., and Tucker, R. S. Fog computing may help to save energy in cloud computing. IEEE Journal on Selected Areas in Communications, 34, 2016.

12. Kumar, P., and Bhatt, A. K. Enhancing multi-tenancy security in cloud computing using hybrid ECC-based data encryption approach. IET Communications, 14, 2020.

13. Li, J., Zhang, Y., Chen, X., and Xiang, Y. Secure attribute-based data sharing for resource-limited users in cloud computing. Computers & Security, 72, 2018.

14. Marinescu, D. C. Cloud Computing: Theory and Practice. 2013.

15. Swami, R., and Das, P. An effective secure data retrieval approach using trust evaluation: HBSEE-CBC. International Journal of Information and Communication Technology, 17, 2020.

16. Varghese, B., and Buyya, R. Next generation cloud computing: New trends and research directions. Future Generation Computer Systems, 79, 2018.

17. Vengala, D. V. K., Kavitha, D., and Kumar, A. S. Secure data transmission on a distributed cloud server using optimized CP-ABE-ECC. Cluster Computing, 23, 2020.

18. Yang, Z., Hu, S., and Chen, K. FPGA-based hardware accelerator of homomorphic encryption for efficient federated learning. Master's Thesis, Hong Kong University of Science and Technology, 2020.

19. Yarava, R. K., and Singh, R. P. Efficient and secure cloud storage auditing based on the Diffie-Hellman key exchange. International Journal of Intelligent Engineering and Systems, 12, 2019.

20. Yu, F., Zhang, H., Zhao, B., Wang, J., Zhang, L., Yan, F., and Chen, Z. A formal analysis of TPM 2.0 HMAC authorization under digital rights management scenario. Security and Communication Networks, 9, 2016.

21. Zha, Y., and Li, J. Virtualizing FPGAs in the cloud. ASPLOS '20, 2020.

22. Zha, Y., and Li, J. Hetero-ViTAL: A virtualization stack for heterogeneous FPGA clusters. ISCA '21, 2021.

23. Zhang, F., Chen, J., Chen, H., and Zang, B. CloudVisor: Retrofitting protection of virtual machines in multi-tenant clouds. SOSP '11, 2011.

24. Zhang, J., Lin, Y., Lyu, Y., and Qu, G. A PUF-FSM binding scheme for FPGA IP protection and pay-per-device licensing. IEEE Transactions on Information Forensics and Security, 2015.

25. Zhang, J., and Qu, G. Recent attacks and defenses on FPGA-based systems. ACM Transactions on Reconfigurable Technology and Systems, 2019.

26. Zhang, F., Wang, Z., Shen, H., Yang, B., Wu, Q., and Ren, K. DARPT: Defense against remote physical attack based on TDC in multi-tenant scenario. DAC '22, 2022.

27. Zeitouni, S., Dessouky, G., and Sadeghi, A.-R. SoK: Security challenges and risks of multi-tenant FPGAs in the cloud. arXiv, 2020.