Chapter 20
Hardware Acceleration

Dina Mahmoud

20.1 Introduction

With Moore's law and Dennard's scaling no longer fueling the improvement in computing performance, new avenues for increasing performance are needed. Hardware acceleration is one avenue where many researchers and industrial parties work and invest. This is because accelerators can allow for high levels of parallelism not supported by general-purpose central processing units. These high levels of parallelism are particularly well-suited for many modern applications. Therefore, research on and use of hardware acceleration is expected to continue soon. However, various parties should consider various aspects when deciding whether to invest in hardware acceleration by making their accelerators or buying them from a third party. This factsheet presents an analysis of hardware acceleration and the trends until 2025. It also discusses the aspects to consider and how specific considerations are more important for some actors.

20.2 Analysis

With the slowdown of Moore's law, system developers are examining potential avenues for performance improvement of computing systems. As simply increasing the frequency or the number of transistors on the chip is no longer feasible, IT infrastructure operators have adopted hardware acceleration. Various platforms and levels of hardware acceleration exist.

D. Mahmoud (✉)
EPFL, Lausanne, Switzerland
e-mail: dina.mahmoud@epfl.ch

V. Mulder et al. (eds.), *Trends in Data Protection and Encryption Technologies*,
https://doi.org/10.1007/978-3-031-33386-6_20

20.2.1 Definition

Hardware acceleration is the use of specialized hardware within a computing system designed to handle specific tasks in an optimized way [1]. Central processing units (CPUs) are typically responsible for most tasks within a computing system. However, tasks requiring high levels of parallelism do not run efficiently on general-purpose CPUs. Moreover, many tasks need to be run, and if one of them is slower than the rest, it can affect the system's performance. This is where hardware acceleration comes in. When a task possesses properties making its execution on a CPU suboptimal, system designers include specialized hardware in the system to which the task is offloaded. Graphics processing units (GPUs) are among the most famous hardware accelerators designed for rendering graphics. Their support for parallelism has also made them suitable for other tasks, including machine learning acceleration [2]. To avoid long execution times due to the sequential nature of CPUs and to avoid software-based exploits, cryptographic algorithms are also among the notable applications benefiting from hardware acceleration [3]. Other examples of hardware accelerators include application-specific integrated circuits (ASICs), which can implement specialized cryptographic accelerators on modern systems-on-chip (SoCs), and field-programmable gate arrays (FPGAs).

20.2.2 Trends

There has been a steady growth in research on hardware acceleration (as shown in Fig. 20.1) and in the adoption of specialized hardware. For instance, specialized accelerators like the Apple Neural Engine are making their way into consumer

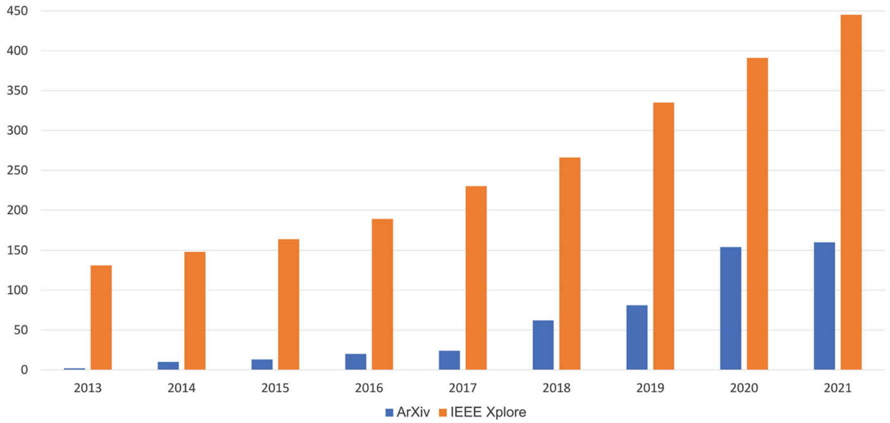

Fig. 20.1 Describes arXiv and IEEE Xplore publications containing the keywords "Hardware acceleration" OR "Hardware accelerator"

electronics [4]. ARM also offers security algorithm accelerators to support the Armv8-A cryptography extensions [5]. Due to the inability of CPUs to meet the increasing computing demand, the use of specialized hardware is expected to keep increasing until 2025. The increased interest in using hardware accelerators has also increased the research on their security in terms of attacks and countermeasures.

20.3 Consequences for Switzerland

Hardware acceleration is helpful for more efficient computing. Nevertheless, there are many considerations when investing in a hardware accelerator, especially when using it for security. If the specialized computing core is to be highly utilized, it is helpful to invest in it [2]. This is a likely case for a cryptographic accelerator in a system that encrypts and decrypts all outgoing and incoming data, for example. However, one needs to consider that the security of hardware accelerators may be questionable. Hardware Trojans, fault injection attacks, and side-channel attacks are significant threats to hardware cryptographic accelerators.

20.3.1 Implementation Possibilities: Make or Buy

This section presents the pros and cons of buying or making secure hardware accelerators (Table 20.1).

Depending on the type of hardware accelerator, many risks and opportunities are associated with making or buying it. Making the hardware accelerator from scratch gives higher security guarantees as there is no possibility for third-party-implanted hardware Trojans. Furthermore, specific countermeasures can be implemented to protect against various exploits, such as redundancy, masking, and hiding. However, there is significant engineering effort in building a correct, highly performant, secure hardware accelerator. If not correctly designed and built, bugs can result in misbehavior, reducing the system's performance or reliability. Furthermore, the interoperability of hardware accelerators with other components in the system is a critical aspect to consider. Buying accelerators usually guarantees that they will have standard interfaces, but it is possible to design a specialized core with standard interfaces.

For specialized military applications, where security is of the utmost concern, and if accelerators are unlikely to already be in existence, making the accelerators is better. Making the hardware eliminates the possibility of hardware Trojans, but not of bugs in the design nor potential leakage of information. Coupled with proper testing (for security and reliability), making the accelerator would guarantee secure and reliable hardware. If the accelerator can be bought, design effort can be saved, but guaranteeing the security would come at the cost of extensive security testing to guarantee that there are no (intentional or unintentional) backdoors. Making

Table 20.1 Implementation possibilities for different sectors

	Make		**Buy**	
	Pros	Cons	Pros	Cons
Military	No hardware Trojans/More specialized hardware	Potential error in implementation adding attack surface	Reuse available tested accelerators, no design effort and accelerator interoperable with other platforms	Might contain accidental or purposeful backdoors and increased need for security testing
Civil Society	None	None	Accelerator interoperable with other platforms, reuse available tested accelerators	Potential security vulnerabilities
Economy	Sell for profit	Liability in case the accelerator has an error or a security vulnerability	Faster development of products	Less advantage over competition features and extensions are needed to innovate

accelerators that target widely used applications can have a significant economic benefit for businesses. Such accelerators can be sold to many parties resulting in high profits. For example, many hardware accelerator designs can be bought and used on Amazon Web Services Marketplace [6]. However, this is only the case for widely used applications. For more specialized accelerators, making them may still prove helpful to the business if the accelerator significantly improves their workloads' performance. However, there may be no direct profit from selling the accelerators. Buying pre-existing accelerators will allow for faster end-product development if the business entity is not accustomed to building hardware. For the remaining actors, buying the accelerators is a good solution. Again, some testing would need to be done to guarantee a minimum level of security for the purchased hardware.

20.3.2 Variations and Recommendation

Hardware accelerators have varying levels of specialization (and flexibility). They can also be integrated using a variety of methods in existing systems. The choice of which variation to opt for depends on the application and the actor looking to use the specialized hardware. The highest level of specialization is achievable when using application-specific integrated circuits (ASICs). However, this translates to higher costs in engineering efforts and reduced flexibility. Field programmable gate arrays (FPGAs) require less effort to design the accelerator and offer more flexibility at the price of remaining within the constraints of the FPGA resources and slightly reduced performance. Finally, graphics processing units (GPUs) offer high flexibility and parallelism. However, they are not as customizable to the application as FPGAs and ASICs and can therefore have lower performance. If the application for which the custom hardware is being purchased will be fully utilizing the hardware, and one in which performance is of the utmost importance, then an ASIC is the best choice. The lower the utilization is expected to be, the better it is to opt for a more flexible option.

Deploying any chosen accelerator still requires considering its security. With attacks constantly being demonstrated against ASIC-, FPGA-, and GPU-based accelerators, designs should be appropriately secured before deployment. The main security risks arise if the device is physically accessible to a remote party. However, software access can also be leveraged for a variety of exploits. According to the deployment model of the device and the desired security level, various protection mechanisms (e.g., redundancy, radiation-hardening, leakage detection) can be implemented.

20.4 Conclusion

The use of specialized hardware to accelerate applications not performing well on modern CPUs will likely continue in the coming years. Consequently, they are likely to be used by all actors in various applications. For example, we already see many hardware accelerators for cryptographic applications and security. Each actor needs to decide on the variation of customized hardware to invest in based on the expected usage and the application. Furthermore, security should be essential when designing or buying the hardware accelerator. The tradeoff between security and design cost must be studied to decide whether to make or buy the accelerator and the amount of testing necessary to guarantee the desired level of security and reliability.

References

1. Wen-mei Hwu and Sanjay Patel. Accelerator Architectures – A Ten-Year Retrospective. *IEEE Micro*, 38(6):56–62, November 2018. Conference Name: IEEE Micro.
2. IRDS$^{\mathrm{TM}}$ 2021: Executive Summary - IEEE IRDS$^{\mathrm{TM}}$. https://irds.ieee.org/editions/2021/executive-summary, August 2022.
3. Mohamed Gafsi, Mohamed Ali Hajjaji, Jihene Malek, and Abdellatif Mtibaa. FPGA hardware acceleration of an improved chaos-based cryptosystem for real-time image encryption and decryption. *Journal of Ambient Intelligence and Humanized Computing*, October 2021.

4. Apple unleashes M1, November 2020. Apple Newsroom.
5. Arm Ltd. Security algorithm accelerators. https://developer.arm.com/downloads/-/security-algorithm-accelerators.
6. AWS Marketplace: Homepage. https://aws.amazon.com/marketplace, August 2022.