

Secure Video Streaming Using Dedicated Hardware

Nicholas Murray-Hill¹, Laura Fontes^{}¹, Pedro Machado^{}^{1*}
and Isibor Kennedy Ihianle^{}¹

^{1*}Department of Computer Science, School of Science and Technology, Nottingham Trent University, ERD220, Clifton Campus, Nottingham, NG11 8NS, Nottinghamshire, UK.

*Corresponding author(s). E-mail(s): pedro.machado@ntu.ac.uk;
Contributing authors: n0821510@my.ntu.ac.uk;
n1119003@my.ntu.ac.uk; isibor.ihianle@ntu.ac.uk;

Abstract

The purpose of this article is to present a system that enhances the security, efficiency, and reconfigurability of an Internet-of-Things (IoT) system used for surveillance and monitoring. A Multi-Processor System-On-Chip (MPSoC) composed of Central Processor Unit (CPU) and Field-Programmable Gate Array (FPGA) is proposed for increasing the security and the frame rate of a smart IoT edge device. The private encryption key is safely embedded in the FPGA unit to avoid being exposed in the Random Access Memory (RAM). This allows the edge device to securely store and authenticate the key, protecting the data transmitted from the same Integrated Circuit (IC). Additionally, the edge device can simultaneously publish and route a camera stream using a lightweight communication protocol, achieving a frame rate of 14 frames per Second (fps). The performance of the MPSoC is compared to a NVIDIA Jetson Nano (NJN) and a Raspberry Pi 4 (RPI4) and it is found that the RPI4 is the most cost-effective solution but with lower frame rate, the NJN is the fastest because it can achieve higher frame-rate but it is not secure, and the MPSoC is the optimal solution because it offers a balanced frame rate and it is secure because it never exposes the secure key into the memory. The proposed system successfully addresses the challenges of security, scalability, and efficiency in an IoT system used for surveillance and monitoring. The Rivest–Shamir–Adleman (RSA) encryption key is securely stored and authenticated, and the edge

device is able to simultaneously publish and route a camera stream feed high-definition images at 14 fps. The proposed system enhances the security, efficiency, and reconfigurability of an IoT system used for surveillance and monitoring. The RSA encryption key is effectively protected by the implementation of an MPSoC with Programmable Logic (PL), which also allows for faster processing and data transmission.

Keywords: IoT, FPGA, Graphics Processor Unit (GPU), PYNQ, MPSoC, security streaming, edge computing

1 Introduction

Internet-of-Things (IoT) has become a central focus in the technology industry, with the development of smart devices that collect and exchange data over the internet. These devices, including sensors, actuators, and other IoT-enabled technologies, are used for a variety of purposes such as analysis, processing, and automation. It is estimated that there will be approximately 75 billion IoT devices in use by 2025 [1]. The use of IoT technology has expanded into various fields including smart energy, industrial factories, transportation, and home automation.

However, the widespread integration of IoT systems in our daily lives also leads to an increase in the amount of data being collected and exchanged over the internet, raising concerns about scalability and the protection of sensitive and private data. Many IoT systems rely on centralised architectures, which process and perform security operations on the cloud [2]. These architectures have limitations such as scalability, transaction speeds, interoperability, and privacy/security [3], and may also be a single point of failure in the event of a breach. Centralised servers that manage keys and act as a single trust authority can pose a significant security risk to other devices on the network, as these keys often form the foundation of security systems, cryptography algorithms, and device authentication/verification [2].

To address these issues, distributed infrastructures have emerged, allowing modern IoT edge devices to perform their own processing and transmit data at the edge. These devices feature robust security mechanisms for secure authentication and secret key storage, which can be used for encryption in a secure, hardware-enforced manner to improve end-device security for IoT systems and data integrity [4]. Field-Programmable Gate Array (FPGA) offer such security features through the use of programmable logic and reconfigurability after manufacturing. This reconfigurability allows devices to be updated in order to keep up with the constantly evolving technology landscape and emerging security threats.

The main contribution of this article is the development and implementation of a secure, efficient, and reconfigurable surveillance and monitoring IoT

system using dedicated hardware. By utilising the Multi-Processor System-On-Chip (MPSoC)'s Programmable Logic (PL) to securely store and authenticate a symmetric key, the proposed system improves the security and integrity of data transmitted from an edge device. Additionally, the use of MPSoCs allow the edge device to simultaneously publish and route a camera stream using a lightweight communication protocol, achieving a high capture rate. The proposed system addresses many of the challenges faced by current IoT systems, including scalability, security, and efficiency.

The remainder of this article is organised as follows: relevant literature and related works to the proposed IoT system proposed in this article is reviewed in Section 2, the methodology used to develop and implement the proposed system is described in Section 3, the results and analysis of the proposed system are presented in Section 4, and finally, the conclusion and future work are discussed in Section 5.

2 Related Work

The IoT is a technology that aims to provide an infrastructure for applications that can coordinate the interaction of people, things, and systems for a specific purpose [5]. These applications do not necessarily have a universally adopted standard, but the architectural model of an IoT system typically consists of three main layers: the perception layer, which uses sensors and microcontrollers to perceive the physical environment; the communication or network layer, which processes and transports data; and the application layer, which uses the data to deliver application-specific services to the user [1]. The communication/network layer uses various technologies to package and transmit data due to the processing and bandwidth restrictions of many IoT devices. Hypertext Transfer Protocol (HTTP), which is commonly used for communication between devices on the internet, is not suitable for low-powered IoT devices due to its fully connection-oriented architecture, large header size, and latency [6, 7]. Established communication protocols that are more suitable for these power and bandwidth-restricted IoT requirements include Constrained Application Protocol (CoAP), Message Queueing Telemetry Transport (MQTT), and Extensive Messaging and Presence Protocol (XMPP). Corak et al. [8] evaluated and compared the performance of these protocols in a real-world IoT testbed. The metrics considered were packet creation time and packet delivery speed to determine the delay differences. The study found that XMPP had the worst performance due to its use of Extensible Markup Language (XML) format, which increased latency. MQTT and CoAP had similar overall performance in terms of packet creation and transmission time, but MQTT was found to be more optimised and standardised. In addition to these protocols, wireless technologies such as Low Range (LoRa), Low Range Wide Area Network (LoRaWAN), and Low-Power Wide Area Network (LPWAN) can be used to enable long-range and low-power communications for IoT devices [9].

These wireless protocols are designed to provide low-power, wide-area networks, making them ideal for use cases where devices need to transmit small amounts of data over long distances, such as in agriculture, smart cities, and industrial applications [9]. Van der Westhuizen and Hancke [10] conducted a more in-depth comparison between CoAP and MQTT to determine which was the most suitable for use with constrained devices, specifically sensors. The comparison considered communication delay and network traffic. Both protocols were found to be good choices for resource-constrained devices, with similar performance and response times. However, the most suitable protocol depended on the overall requirements of the system. CoAP was found to be the optimal choice for interfacing with business systems, due to its small average packet sizes and minimal battery/data usage. MQTT, on the other hand, was found to be the preferred solution for systems such as home automation and sensor networks, where device heterogeneity is more pronounced. MQTT was easier to configure for new devices and had the most effective data flow thanks to its publish/subscribe model and use of Quality of Service (QoS).

2.1 Edge computing

Edge computing is a network architecture that involves processing sensory (e.g. visual data) data closer to the source, rather than on the cloud [11]. This allows for fast processing and efficient handling of data intensive operations in real-world scenarios such as the IoT. While edge computing can offer benefits for IoT systems, there are also limitations in terms of security. Khan et al. [11, 12] found that further development is needed in areas such as authentication and access control, and that tamper-proof architectures may be one solution to addressing these security issues. However, securing large scale and time-critical IoT systems can also be challenging due to the cost of methods such as encryption in terms of latency, energy consumption, and network bandwidth [13]. Additionally, the heterogeneity of devices that communicate across these networks without a well-established protocol can also pose challenges [1]. Fortunately, professionals in the field are working to overcome these limitations and improve the safety and efficiency of communication between IoT devices.

2.2 FPGA technology

FPGAs are specialised hardware devices that have gained popularity in the edge computing space due to their ability to solve problems through reconfigurable hardware circuits. These circuits can be described using Hardware Description Languages (HDL) such as Verilog and Very High Speed Integrated Circuit Hardware Description Language (VHDL), and are made up of various logic units such as look-up tables, flip-flops, and multiplexers. FPGAs offer several benefits for security, parallel computing, and flexibility to update hardware designs after deployment [14–16]. They have also been advanced through the use of System-on-Chip (SoC), which integrate programmable logic with


real-time processors. An example of this is the AMD-Xilinx Zynq Ultrascale+ MPSoC¹, which includes an Advanced Reduced Instruction Set Computing Machine (ARM) CPU, programmable logic, and units for graphics and video processing. While FPGAs and SoCs have similarities with microcontrollers, FPGAs offer advantages in physical and cybersecurity through encrypted bit-streams and key loading mechanisms, and can act as a Root of Trust (RoT) by holding security private keys and critical algorithms. FPGAs also show greater efficiency in processing algorithms for image processing and video transcoding due to their parallel computing capabilities.

While FPGAs offer significant advantages for IoT, they are considered complex due to the low-level hardware knowledge required, such as VHDL and Verilog. To address this, FPGA vendors have been promoting the use of high-level design flows and tools that allow for the creation of Register Transfer Level (RTL) designs using high-level languages like C, C++, System C and Open Computer Language (OpenCL). However, the question remains as to how well these high-level designs compare to manually written RTL designs in terms of optimisation. Guo et al. [17] discussed that while High Level Synthesis (HLS) may not be as optimised as manually written RTL designs for complex designs, the use of directives like loop unrolling and loop merging and pipelining can significantly improve resource utilisation, reduce latency, increase resource sharing, and optimize logic for video processing algorithms. These findings suggest that FPGA technology can be more accessible to designers without strong low-level hardware knowledge, while still maintaining good performance.

In summary, the reviewed articles have demonstrated the various considerations and challenges faced in the design and implementation of an IoT system. Communication protocols such as MQTT and CoAP have been shown to be effective in resource constrained environments, but the choice between them ultimately depends on the specific requirements of the system. Edge computing has the potential to improve the efficiency and security of IoT networks, but also comes with its own limitations that require further development. FPGA technology offers advanced security and parallel processing capabilities for IoT, but can be complex to implement. High level synthesis tools, such as the AMD-Xilinx Vivado HLS², have been shown to improve the productivity and performance of FPGA designs for real time image processing applications, but may not always be as optimised as manually written designs. These findings highlight the importance of carefully evaluating the various technologies and approaches available for a particular IoT system in order to ensure optimal performance and security.

¹ Available online, <https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html>, last accessed 07/01/2023

² Available online, <https://www.xilinx.com/support/documentation-navigation/design-hubs/dh0090-vitis-hls-hub.html>, last accessed 07/01/2023

**Fig. 1** Architecture Diagram

3 Methods

The proposed IoT system utilises the Ultra96-V2 Development Board (Ultra96) equipped with a powerful AMD-Xilinx Zynq UltraScale+ MPSoC ZU3EG³ device as the main processing system at the perception layer. The performance of the Ultra96-V2 Development Board (Ultra96) was compared to a NVIDIA Jetson Nano (NJN) and a Raspberry Pi 4 (RPI4) under the same testing conditions.

To establish a fair comparison, each processing device (i.e. Ultra96, RPI4 and NJN) runs an MQTT client to publish data from its connected Universal Serial Bus (USB) webcam to an MQTT broker, which acts as an intermediary to route the data to interested parties. The camera feed is then displayed on a Node-RED⁴ dashboard at the application layer for subscribers to view. The use of MQTT and the Node-RED dashboard allows for efficient and flexible communication and data management within the system. The system also implements security measures, such as bitstream authentication, to protect against potential attacks. Overall, the proposed IoT system utilises a variety of technologies to coordinate the interaction of people, things, and systems for a specific purpose (see Figure 1).

The Avnet Ultra96, which is powered by an AMD-Xilinx Zynq UltraScale+ MPSoC device that combines an ARM processor and FPGA. The Ultra96 is energy efficient and performs well due to designated processors being responsible for specific tasks. The Processor System (PS) in the AMD-Xilinx Zynq UltraScale+ MPSoC runs an ARM64v8 Linux environment for running a web server while also interfacing with the programmable logic via the ARM eXtensible Interface 4 (AXI4) for user authentication and on-field reconfigurability.

³Available online, <https://www.xilinx.com/content/dam/xilinx/imgs/products/zynq/zynq-eg-block.PNG>, last accessed 07/01/2023

⁴Available online, <https://nodered.org/>, last accessed 07/01/2023

ARM64v8 is a version of the ARM architecture that supports 64-bit instructions. It is used in some 64-bit ARM processors, such as those used in the Avnet Ultra96.

The NJN⁵ is a small, powerful computer designed for use in image and video processing applications. It is powered by a quad-core ARM Cortex-A57 CPU and a 128-core NVIDIA Maxwell GPU, running on an ARM64v8 Linux environment. Programming the NJN is typically done using the NVIDIA JetPack SDK (NJPSDK), which includes a Linux-based development environment and a variety of software libraries for working with the CPU and GPU. Open Computer Vision (OpenCv) compiled with the CUDA library was used to achieve maximum performance when processing visual data.

The RPI4⁶ is a low-cost, single-board computer designed for educational and hobbyist use. It is powered by a Broadcom BCM2711, quad-core Cortex-A72 (ARMv8) 64-bit SoC @ 1.5GHz and runs on a Linux-based operating system, typically Raspbian. OpenCv was used to achieve maximum performance when processing visual data. The RPI4 is known to be cost-effective solution but with lower frame rate compared to other devices.

The IoT system was not initially configured with any software or input sensors, so a USB webcam was connected to the USB 3.0 Type A port to capture the camera feed. The Micro-B upstream port was used to connect to a host workstation on the same local network, although the board also supports WiFi connectivity. The device was booted using a Micro Secure Digital (uSD) card that loaded the Python productivity for Zynq (PYNQ) framework. This open-source framework allows developers to program AMD-Xilinx UltraScale+ Zynq FPGA devices, such as the one used in this device, using Python. Furthermore, PYNQ was designed to be used in embedded systems and provides a set of libraries, drivers and Jupyter notebooks to enable easy programmability of FPGAs through high-level programming languages like Python. This setup allows for the physical connection and control of the camera feed, which can be streamed in the proposed system.

3.1 PYNQ Framework

Bootting the device required software to be loaded onto an SD card, in order to leverage the security and parallel hardware execution benefits of the Ultra92-V2 programmable logic, the approach was to use the PYNQ framework version 2.6⁷. This platform features a Linux operating system, along with the Python software package and a Jupyter⁸ web server for developing solutions on the board for rapid on-field development and reconfigurability over a network. This image should be flashed onto a uSD card with a capacity of at least 16GB and inserted into the board. Once powered on, the board can be accessed

⁵ Available online, <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>, last accessed 15/01/2023

⁶ Available online, <https://www.raspberrypi.com/products/raspberry-pi-4-model-b/>, last accessed 15/01/2023

⁷ Available online, <http://www.pynq.io/>, last accessed: 05/05/2022

⁸ Available online, <https://jupyter.org/>, last accessed: 05/05/2022

by connecting a USB Micro-B cable to a host PC or by setting up a WiFi connection on the local network. The board is configured with the default IP address of 192.168.3.1, which allows access to the locally hosted Jupyter web server.

This revised statement provides more clarity by breaking up the original sentence into several shorter sentences. It also provides more detail on how to access the board, including both USB and WiFi options, and clarifies that the Jupyter web server is hosted locally. Additionally, it uses the passive voice as requested.

3.2 CUDA

To develop and execute various components which run on the board, there are prerequisites that should be installed during the setup phase. OpenCv version 4.5.1⁹ for Python is used to retrieve frames from an input device, such as a webcam or IP camera. NumPy version 1.16.0¹⁰ is also used within the Jupyter notebook to manipulate data structures, such as arrays. This was used to read and write user's credential files. An MQTT broker should also be installed to route the data between publisher and subscribers, so the Mosquitto-MQTT version 1.4.15¹¹ was installed from Ubuntu's open-source universe repository. Finally, to create an MQTT client to publish the stream from the embedded processing system, the Paho-MQTT¹² Python package was installed. At this point, all the prerequisites to develop and execute the system are in place on the PYNQ Linux environment. A detailed list of the tools and equipment used in the system can be found in Table 1.

Table 1: List of equipment and tools

Resource	Type	Description
Anvet Ultra96	Equipment	System on Chip device for authenticating user key and running the various components of project
NJN	Equipment	Embedded system capable of achieving higher frame rates and running the various components of project
RPI4	Equipment	Low cost embedded system, equipped with 2GB of DDR, ideal for reducing the costs and running the various components of project

⁹ Available online, <https://opencv.org/>, last accessed: 05/05/2022

¹⁰ Available online, <https://numpy.org/>, last accessed: 05/05/2022

¹¹ Available online, <https://mosquitto.org/>, last accessed: 05/05/2021

¹² Available online, <https://pypi.org/project/paho-mqtt/>, last accessed: 05/05/2022

Full HD 1080p Web Camera	Equipment	Camera used to capture environment for smart camera
Vivado HLS 2020.1	Development Tool	AMD-Xilinx tool used for creating custom authentication Intellectual Property (IP) core
Vivado 2020.1	Development Tool	AMD-Xilinx tool used for creating hardware design and generating the secure bitstream
PYNQ	Framework	Framework for booting the Ultra96 board with Ubuntu 18.04 and interfacing with the programmable logic
OpenCV	Library	Library used to capture input streams, such as camera and online video stream
NVIDIA CUDA 11	Development Tool	NVIDIA tool used for accelerating the video streaming using CUDA IP cores
NumPy	Library	Various array manipulation functionalities and for saving/reading the user credentials file
Paho MQTT	Library	Used to create an MQTT client on the Ultra96 to publish the camera stream to a topic
Mosquitto-MQTT	Software	Used to run a MQTT broker which listens to client publishing to manage and route to the correct subscribers
Node-RED	Development Tool	The application layer is created using Node-RED to create a IoT flow
PuTTY	Terminal Emulator	An emulator which has built in protocols such as SSH to connect to the Ultra96

3.3 Designing the Secure Bitstream of the Ultra96

To secure the system, a bitstream file was created to provide confidentiality through a 256-bit secret key and method of authentication. The key is described in the FPGA logic and is embedded within the FPGA unit, allowing the system to securely store the private key and prevent it from being exposed in the RAM. The system employs a pair of keys, consisting of a public and a private key, to ensure secure message authentication. Nevertheless, the private key is securely stored in the FPGA logic, making the proposed method safer

than other authentication methods that store private keys externally. The use of a private key stored in the FPGA logic ensures that only authorised devices with the corresponding public key are granted access to the system and camera stream. This is achieved through a secure authentication process where the authorised device sends an authentication message encrypted with the public key. The FPGA then decrypts the message using the private key stored in its logic to confirm that the device is authorised for providing an additional layer of security to the system, making it more resistant to unauthorised accesses. The high level design flow was used to develop the hardware design at a higher level of abstraction using C/C++ code, which was then converted into optimised RTL code by a compiler. Custom Intellectual Property cores were also included in the design and interfaced with via the PYNQ framework to run on the Ultra96 programmable logic.

The process of building the authentication IP core begins with using AMD-Xilinx Vivado HLS¹³ software. Using the HLS software, the top level function was written in C containing a secret key, authentication method to compare the valid key with the input key, along with the required I/O ports for the PL to interface with the IP block. In this case there were two ports, one of which was the key with a size of 256-bits and the other being the authentication result which was a single bit boolean value to represent whether the input key was valid or not. Due to the size of these ports being relatively small, the AXI4-Lite protocol was utilised for the Ultra96 processing system to interface with the IP block, as this is generally a suitable design choice for smaller data transfers. To optimise the design, various loop and array optimisation directives were tested in the hope to reduce the estimated clock time and maximum clock cycles. By using the pipeline *Pragma* in the loop to compare the keys, the maximum clock cycles was reduced from 64 to 34. This directive works by reducing the initiation interval for the loop by allowing concurrent execution of the operations. To verify the output of the top level function a test bench was written, this test bench was used by the HLS tool during C simulation, synthesis and C/RTL co-simulation to validate that the produced RTL was functionally identical to the C code that was written and therefore confirming that the IP is working as intended to be packed and exported. The timing and latency summary for the IP is presented in Figure 2.

To verify the functionality of the custom IP block, a test bench was written in C. This test bench was used by the HLS tool during C simulation, synthesis, and C/RTL co-simulation to validate that the produced RTL was functionally identical to the C code, and therefore confirm that the IP was working as intended. Once the custom IP block was exported, it could be used as part of the wider system by importing it into the AMD-Xilinx Vivado Design Suite. This tool has an IP Integrator, which was used to build the hardware design by integrating the custom IP block with IPs available in the AMD-Xilinx's IP catalogue. A block diagram, shown in Figure 3, was generated using the Vivado

¹³Available online, <https://www.xilinx.com/products/design-tools/vivado.html>, last accessed: 05/05/2022

The IP block is lightweight and only uses a small portion of the programmable logic resources on the device. Figure 4 shows the resource utilisation.

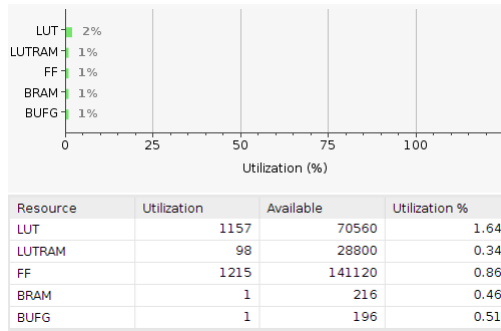


Fig. 4 IP Resource utilisation

PYNQ is accessed through a local Jupyter web server at 192.168.3.1. It allows the execution of Python packages and libraries on the Ultra96 board. The PYNQ Overlay class can be used to view and interface with the PL of the Ultra96 using the previously created bitstream and default overlay driver to access the IP's ports configured in the drivers file. The authentication result is retrieved using three specific addresses: the start control signal (0x000), the offset of the input key port (0x080), and the offset of the data out port (0x100). The user's symmetric key can be loaded into the input key port a 4-byte integer at a time. The start control signal is set to high to start the IP and the authentication output is read from the data out port. If the key is valid, the camera is initialised and published to the MQTT broker.

The MQTT broker was configured to automatically start on boot and run on localhost with the port 1883. OpenCv and Paho-MQTT were also imported for use in capturing and publishing the camera stream. A configuration file was placed on the device to define system parameters such as camera settings, MQTT settings, and the path to the credentials file. This file allows the user to easily update system parameters without requiring knowledge of the system. The data visualisation at the application layer was the final component for the project. This layer is responsible for connecting the clients or subscribers to the MQTT broker and displaying the secured camera feed. Node-RED is a browser-based editor, where flows can be built using a catalogue of nodes to fit custom IoT requirements. Additional nodes can also be installed via node package manager. This tool was chosen for the project due to being open source and high productivity, where additional nodes can be quickly inserted into the flow and deployed instantly. The flow that was designed consisted of an MQTT input node, which is configured to the MQTT broker running on the host workstation. This node receives the message payload from the broker as a base64 string and then passes this into an HTML image template, which is finally connected to a dashboard widget template where it is displayed automatically. Once this flow is deployed, the dashboard can be accessed on the local network at the URL 127.0.0.1:1880/ui. The designed flow is shown in Figure 5.



Fig. 5 Node-RED Flow. The streaming node that will be enabled only after a successful authentication process.

4 Results

The aim of this project was to build a flexible and reconfigurable edge device that could protect the integrity of data within a surveillance and monitoring IoT system. To achieve this, a secure authentication mechanism was implemented to guarantee that the edge device could only publish the camera stream when data integrity and authenticity could be assured. This was accomplished by concealing a 256-bit secret key and method of authentication inside a bitstream file, which is the hardware description of an FPGA and is difficult to reverse engineer due to being a stream of bits that only describe the hardware logic itself. This provided the necessary confidentiality to protect the key.

The proposed system was tested under the same conditions on the Ultra96, NJN and RPI4 to ensure that each version of the IoT device delivered the expected functionality and behaviour. Several tests focused on the integration of the IP within the overall system were carried out to ensure that the system is working as intended. This involved testing the end-to-end process of capturing the camera stream, publishing the data over MQTT, and displaying the stream on the dashboard. These tests were carried out by setting up the Ultra96 board, NJN and RPI4 with the boot image and prerequisites, importing the bitstream, and running the python script to capture and publish the camera data. The Node-RED flow was then set up, and the dashboard accessed to verify that the stream was being displayed as expected on all devices. Additionally, the system was tested under various scenarios such as using the correct key, using an incorrect key, and attempting to access the stream without providing a key, to ensure that the system was functioning as intended and that the authorisation process (see Figure 7). was secure (see Figure 7). These tests were important to ensure that the Ultra96, NJN and RPI4 provide optimal solution, balanced frame rate and secure key storage. It is clear from Table 2 that the hardware implementation (i.e. Ultra96) had the same performance has the software implementation (i.e. NJN and RPI4) but without exposing the private key into the RAM.

Ultra96, the top level function was written in C, containing the secret key and authentication method to compare the valid key with the input key, along with the required I/O ports for the PS to interface with the IP block. The AXI4-Lite protocol was used for the Ultra96 processing system to interface with the IP block, as it is suitable for smaller data transfers. To optimize the design, various loop and array optimisation directives were tested to reduce the

14 *Secure Video Streaming Using Dedicated Hardware*

```

4Running IP key tests
5
6Testing Correct Key: 7336763979244226452948404d635166546a576e5a7234753777217a25432a46
7Success: 1
8
9Testing Invalid Key: 472b5a391a244226452948144d635166546a576e5a72347537773b7a3f432a1e
10Success: 1
11
12Testing Invalid Key: 4d2b4a391a244226452948144d2a51662a6a3b6e1f72347537773b7a3f432a1c
13Success: 1
14
15Testing Invalid Key: 7436763979244226452948404d635166546a576e5a7234753777217a25432a46
16Success: 1
17
18Testing Incomplete Key: 7336763979244226452948404d635166546a576e5a7234753777217a25432a0
19Success: 1
20
21Testing Incomplete Key: 36763979244226452948404d635166546a576e5a7234753777217a25432a60
22Success: 1
23
24Testing Empty Key
25Success: 1
26
27Tests passed: 7/7
28INFO: [SIM 1] CSim done with 0 errors.
29INFO: [SIM 3] ***** CSIM finish *****

```

Fig. 6 Authentication Unit Tests. The output of the seven unit tests, where the first test involved providing the correct key followed by three tests with incorrect keys, and the last three tests involved two incomplete keys.

estimated clock time and maximum clock cycles. By using the pipeline *Pragma* in the loop to compare the keys, the maximum clock cycles was reduced from 64 to 34.

Table 2 Authentication process (number of attempts) per each processing system

Tests	RPI4	NJN	Ultra96
Correct Key	10	10	10
Invalid Key	5	5	5
Incomplete Key	7	7	7
Empty Key	4	4	4
Wrong key	6	6	6
Successful authentications	10	10	10
Unsuccessful authentications	22	22	22

To verify the output of the top level function, in the Ultra96, a test bench was written and used by the HLS tool during C simulation, synthesis, and C/RTL co-simulation to validate that the produced RTL was functionally identical to the C code and that the IP was working as intended. Once the custom IP block was exported, it could be used as part of the wider system by importing it into the AMD-Xilinx Vivado Design Suite. This tool has an IP Integrator, which was used to build the hardware design by integrating the created IP block with IPs from the AMD-Xilinx's IP catalogue. A single memory-mapped AXI master and AXI slave Interconnect was included to ensure that the system was able to handle various scenarios that may occur during operation. These tests included different combinations of correct and

incorrect keys, as well as edge cases such as missing or incorrect bytes in the key. It was essential that all of these tests passed in order to consider the IP secure and fit for purpose. In addition to these unit tests, the overall system was also tested to ensure that it functioned as intended. This included testing the MQTT communication protocol, the PYNQ framework, and the Node-RED dashboard. Overall, the testing of the system showed that the objective of building a flexible and reconfigurable edge device was met, as the system was able to securely store and use a secret key and authentication method, and was also easily configurable and adjustable through the use of a configuration file and various software tools. A demo video [18] is available on YouTube¹⁴ demonstrating the system working.

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1458	root	20	0	1067016	102452	39112	S	41.4	5.0	0:20.70	python3
102	root	0	-20	0	0	0	R	13.9	0.0	0:09.10	kworker/u9+
653	root	20	0	211180	54452	13240	S	3.6	2.7	0:12.93	jupyter-no+
402	root	-51	0	0	0	0	S	1.7	0.0	0:01.50	irq/46-dwc3
426	mosquit+	20	0	7512	4864	4136	S	1.3	0.2	0:01.23	mosquitto
9	root	20	0	0	0	0	S	0.7	0.0	0:00.63	ksoftirqd/0
1478	xilinx	20	0	8700	3132	2572	R	0.7	0.2	0:00.10	top
293	root	20	0	0	0	0	S	0.3	0.0	0:00.04	kpktgend_3
1	root	20	0	13944	8428	5672	S	0.0	0.4	0:02.91	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd
3	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	rcu_gp
4	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	rcu_par_gp
7	root	20	0	0	0	0	I	0.0	0.0	0:00.02	kworker/u8+
8	root	0	-20	0	0	0	I	0.0	0.0	0:00.00	mm_percpu_+
10	root	20	0	0	0	0	I	0.0	0.0	0:00.14	rcu_sched
11	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	migration/0
12	root	20	0	0	0	0	S	0.0	0.0	0:00.00	cpuhp/0

Fig. 7 Ultra96 CPU resource consumption

The final results obtained for all the devices are listed in Table 3.

Table 3 Results comparison

Device	Security aspects	Frame rate no auth	Frame rate with auth	Image size	Average price
Ultra96	Safest because the private key is never exposed to memory.	14 near real-time	14 near real-time	1920 × 1080	£300.29 GBP ¹⁵
RPI4	Unsafe because the secure key was exposed to the memory.	6 far real-time	7 far real-time	1920 × 1080	£52.00 GBP ¹⁶
NJN	Unsafe because the secure key was exposed to the memory.	30 real-time	30 real-time	1920 × 1080	£61.00 GBP ¹⁷

The proposed system was tested under the same conditions on the Ultra96, NJN and RPI4 to evaluate the performance and security of the IoT device. The results showed that the NJN achieved the highest frame rate of 30 fps (real-time), making it the best in terms of frame rate. The RPI4 offered a more cost-effective solution but with a lower frame rate of 6 fps, making it the worst in terms of frame rate. And the Ultra96 achieved a frame rate of 14 fps and offered a safer solution by securely storing the RSA encryption key in the FPGA unit, making it the best in terms of security and performance.

¹⁴ Available online, <https://youtu.be/8AXlf6tRZyo>, last accessed 07/01/2023

5 Discussion and Future Work

In comparison to other authentication systems that rely solely on the use of a CPU, the proposed IoT system utilizing an FPGA has several advantages. One main advantage is the improved security provided by storing the secret key and authentication method within the FPGA bitstream, as it is not accessible in a readable format outside the device. This is in contrast to a CPU-only system where the secret key and authentication method may be stored in plaintext or encrypted in memory, which could potentially be accessed by an attacker with the appropriate tools and knowledge. On the CPU side, private keys are often stored in the file `$HOME/.ssh/id_rsa` as plain text, which poses a security risk. However, tools like `valgrind`¹⁸ and Hex editors can be utilised to inspect processes loaded into memory, making it possible to detect any potential breaches. Additionally, one can enhance security by adding authorised public keys to `$HOME/.ssh/authorized_keys`. Nevertheless, in the case of FPGA logic, keeping authorised public and private keys described within the logic makes it exceptionally challenging to breach security. This is because FPGAs are configured with hardware-level security features that can help to prevent unauthorised access and ensure that the keys remain secure. Although there are risks associated with storing private keys in plain text on the CPU side, there are also measures that can be taken to mitigate these risks, and the use of FPGA logic can provide an additional layer of security.

Another advantage of the proposed system is its reconfigurability and flexibility. With the use of a configuration file, the user is able to easily update the location of the bitstream, as well as change the camera input and MQTT settings without requiring in-depth knowledge of the system. This is not necessarily possible with a CPU-only system, as changes to the system may require modifications to the codebase and potentially require the expertise of a software developer. Therefore, the user can decide which programming system to use based on the budget, security and frame rate restrictions. In terms of performance, the Ultra96 is able to stream the camera feed at a maximum of 14 fps, which is higher than the aim of 6 fps offered by the RPI4. This is due to the efficient resource utilisation of the FPGA, as seen in Figure 7 where the device is only utilizing 63.6% of its CPU resources. In comparison, a CPU-only system may struggle to handle the processing demands of the camera streaming and MQTT communication simultaneously, potentially resulting in lower frame rates or slower performance.

In summary, the proposed IoT system utilising an FPGA for authentication offers improved security, reconfigurability, and performance compared to systems that rely solely on a CPU. There are many directions in which future work on this project could go. One possible avenue of research is to improve the security of the system by implementing more advanced forms of authentication. For example, instead of using a simple symmetric key, a more secure method such as a public-private key pair could be used. This would require

¹⁸Available online, <https://valgrind.org/>, last accessed: 25/03/2023

the use of a cryptographic accelerator or hardware security module to ensure that the key operations can be performed quickly and efficiently on the edge device. Another possibility is to incorporate additional security measures to protect against physical tampering with the device. This could include the use of tamper-evident seals or hardware-based intrusion detection to alert the user if the device has been opened or tampered with.

Another area for improvement could be to optimize the system for better resource utilisation. This could involve using more advanced optimisation techniques during the design phase, or implementing more efficient protocols for communication between the different components of the system. Finally, it would be interesting to explore the possibility of implementing machine learning algorithms on the edge device to enable more advanced forms of data analysis and decision-making. This could involve training a model on the device to identify certain patterns or characteristics in the data, and then using this model to make decisions about how to handle the data. Overall, there are many exciting directions in which this project could be taken, and we believe that it has the potential to make a significant impact in the field of edge computing and IoT security.

Acknowledgements

The authors would like to express their gratitude to Mr Flemming Christensen and Sundance Multiprocessor Technology for their invaluable support and assistance in the form of AMD-Xilinx training.

Declarations

The manuscript has no associated data and the source code can be made available upon request.

References

- [1] Nord, J.H., Koohang, A., Paliszkievicz, J.: The internet of things: Review and theoretical framework. *Expert Systems with Applications* **133**, 97–108 (2019)
- [2] Kouicem, D.E., Bouabdallah, A., Lakhlef, H.: Internet of things security: A top-down survey. *Computer Networks* **141**, 199–221 (2018)
- [3] Farahani, B., Firouzi, F., Luecking, M.: The convergence of iot and distributed ledger technologies (dlt): Opportunities, challenges, and solutions. *Journal of Network and Computer Applications* **177**, 102936 (2021)
- [4] Johnson, A.P., Chakraborty, R.S., Mukhopadhyay, D.: A puf-enabled secure architecture for fpga-based iot applications. *IEEE Transactions on Multi-Scale Computing Systems* **1**(2), 110–122 (2015)
- [5] Wortmann, F., Flüchter, K.: Internet of things. *Business & Information Systems Engineering*, v. 57. Available online, <https://link.springer.com/content/pdf/10.1007/s12599-015-0383-3.pdf>, last accessed: 07/01/2023 (2015)
- [6] Patel, C., Doshi, N.: A novel mqtt security framework in generic iot model. *Procedia Computer Science* **171**, 1399–1408 (2020)
- [7] Shah, S.W.H., Mian, A.N., Mumtaz, S., Wen, M., Hong, T., Kadoch, M.: Protocol stack perspective for low latency and massive connectivity in future cellular networks. In: *ICC 2019-2019 IEEE International Conference on Communications (ICC)*, pp. 1–7 (2019). IEEE
- [8] Çorak, B.H., Okay, F.Y., Güzel, M., Murt, Ş., Ozdemir, S.: Comparative analysis of iot communication protocols. In: *2018 International Symposium on Networks, Computers and Communications (ISNCC)*, pp. 1–6 (2018). IEEE
- [9] Queralta, J.P., Gia, T.N., Zou, Z., Tenhunen, H., Westerlund, T.: Comparative study of lpwan technologies on unlicensed bands for m2m communication in the iot: Beyond lora and lorawan. *Procedia Computer Science* **155**, 343–350 (2019)
- [10] van der Westhuizen, H.W., Hancke, G.P.: Practical comparison between coap and mqtt-sensor to server level. In: *2018 Wireless Advanced (WiAd)*, pp. 1–6 (2018). IEEE
- [11] Khan, W., Ahmed, E., Hakak, S., Yaqoob, I., Ahmed, A.: Edge computing: A survey, future generation computer systems (2019)

- [12] PremSankar, G., Di Francesco, M., Taleb, T.: Edge computing for the internet of things: A case study. *IEEE Internet of Things Journal* **5**(2), 1275–1284 (2018)
- [13] Mohanty, S.N., Ramya, K., Rani, S.S., Gupta, D., Shankar, K., Lakshmanaprabu, S., Khanna, A.: An efficient lightweight integrated blockchain (elib) model for iot security and privacy. *Future Generation Computer Systems* **102**, 1027–1037 (2020)
- [14] Elnawawy, M., Farhan, A., Al Nabulsi, A., Al-Ali, A.-R., Sagahyroon, A.: Role of fpga in internet of things applications. In: 2019 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT), pp. 1–6 (2019). IEEE
- [15] Teubner, J., Woods, L.: Data processing on fpgas. *Synthesis Lectures on Data Management* **5**(2), 1–118 (2013)
- [16] Hoozemans, J., Peltenburg, J., Nonnemacher, F., Hadnagy, A., Al-Ars, Z., Hofstee, H.P.: Fpga acceleration for big data analytics: Challenges and opportunities. *IEEE Circuits and Systems Magazine* **21**(2), 30–47 (2021)
- [17] Guo, K., Zeng, S., Yu, J., Wang, Y., Yang, H.: A survey of fpga-based neural network accelerator. arXiv preprint arXiv:1712.08934 (2017)
- [18] Murray-Hill, N., Machado, P.: Secure Video Streaming Using Dedicated Hardwar. Zenodo. DOI, <https://doi.org/10.5281/zenodo.7513129> (2023). <https://doi.org/10.5281/zenodo.7513129>. <https://doi.org/10.5281/zenodo.7513129>