



# FPGA-Based Acceleration Techniques for High-Throughput Data Processing

Priya Sharma

Data Science Engineer, Cloud Computing Division  
Microsoft Research, India

**Abstract** - Accelerating data processing is crucial across various industries. Field-Programmable Gate Arrays (FPGAs) have emerged as effective tools for speeding up data-intensive processes due to their reconfigurability and parallel processing capabilities. FPGA inline accelerators offer a hardware parallel platform that efficiently handles real-time data analytics workloads. They ingest data while performing inline processing for data conditioning, providing real-time insights from streaming data. By placing the analytic pipeline close to the point of ingress and leveraging hardware acceleration for initial data analysis, FPGAs mitigate latency issues associated with traditional data analytics technologies. FPGAs' inherent architecture allows for massively parallel, real-time processing, making them suitable for early-stage processing tasks such as initial filtering, transformations, data cleaning, and enrichments. FPGA inline accelerators reside between the network interface card (NIC) and the CPU, intercepting incoming packets and enabling functions like transformations, pattern detectors, filters, and compression/decompression in real time. This approach provides invaluable insights in data analytics and reduces data volumes before they pass through software stacks, offloading the CPU. Furthermore, FPGA-based accelerators offer a deterministic approach regardless of data rate or formats, simplifying the system by eliminating the need for complex flow control and load balance management.

**Keywords** - FPGAs, Inline acceleration, Data processing, High-throughput, Real-time analytics, Parallel processing

## 1. Introduction

The escalating volumes of data generated daily across various domains, from scientific research and financial modeling to real-time video processing and network security, demand innovative approaches to data processing. Traditional CPU-based systems often struggle to keep pace with these ever-increasing data rates, leading to performance bottlenecks and increased latency. Consequently, there is a growing need for specialized hardware acceleration techniques that can efficiently handle high-throughput data processing workloads. Field-Programmable Gate Arrays (FPGAs) have emerged as a compelling solution to address these challenges, offering a unique combination of flexibility, performance, and energy efficiency.

### 1.1 The Rise of FPGA Acceleration

FPGAs stand out as reconfigurable hardware devices, allowing developers to customize their architecture to perfectly match the specific requirements of a given application. Unlike CPUs and GPUs, which are designed for general-purpose computing, FPGAs can be programmed at a low level to implement highly parallel and pipelined data processing pipelines. This fine-grained control over the hardware enables the creation of custom accelerators that can significantly outperform traditional processors in data-intensive tasks. The ability to reconfigure FPGAs also provides a distinct advantage, as algorithms evolve, the logic implemented on the FPGA can be updated to reflect these changes.

### 1.2 Advantages of FPGA-Based Solutions

The use of FPGAs for data processing acceleration offers several key advantages. Their inherent parallelism enables massive data throughput, as multiple operations can be performed simultaneously. Furthermore, FPGAs excel in tasks requiring low latency and real-time processing due to their deterministic execution and ability to minimize data movement. This makes them particularly well-suited for applications where timely responses are critical, such as high-frequency trading or network intrusion detection. FPGAs can also be more energy-efficient than CPUs and GPUs for specific workloads.

## 2. Related Work

The use of FPGAs for acceleration has been explored across a variety of computing domains, including image processing and data analytics. This section provides an overview of relevant research in FPGA-based acceleration techniques, highlighting key approaches and their applications.

One prominent area of research involves the acceleration of Convolutional Neural Networks (CNNs) on FPGAs. FPGAs are particularly well-suited for CNN acceleration due to their ability to implement custom hardware configurations tailored to

convolution operations. A comprehensive review by Zhang et al. (2024) covers the history of CNNs, explains key layers, and summarizes FPGA optimization methods by category. Optimizations for both software deployment and hardware design have been implemented to improve computing on FPGAs, unlocking the potential of deploying CNNs on resource-constrained devices. Algorithm optimizations, such as using the Winograd algorithm and Fast Fourier Transform (FFT), can significantly reduce computational complexity and improve computational efficiency. Hardware design optimization also plays a vital role in accelerator FPGAs.

For image processing applications, researchers have developed FPGA-based soft processors like the Image Processing Processor (IPPro), which can operate at high frequencies. Vertical scaling can exploit data parallelism by mapping an actor on multiple processor cores. Image processing hardware acceleration can be achieved using GPUs, TPUs, FPGAs, ASICs, and CPUs. FPGAs allow for custom hardware configurations tailored to convolution operations, making them particularly useful in low-power, real-time applications like embedded systems. Convolution in FPGAs is accelerated through techniques like pipelining, loop unrolling, and dataflow optimization.

In data centers, FPGAs are gaining traction as accelerators to offload compute-intensive tasks from CPUs and GPUs. Microsoft Azure has deployed FPGAs in their data centers to accelerate various services, such as AI inference, data analytics, and networking. FPGA-based acceleration can significantly reduce latency by bypassing the need for data transfers between the CPU and GPU, which is especially beneficial for real-time applications.

Other research has focused on algorithm-hardware co-optimization approaches to achieve CNN acceleration on FPGAs for remote sensing image processing via SIMD architecture. Model optimization techniques can significantly reduce the hardware resource requirements of the model and improve the energy efficiency of the system through operations such as operation fusion and depth-first mapping. The use of Binarized Neural Networks (BNNs), where weights and activation values are restricted to binary values, can also reduce memory utilization.

### 3. Proposed Methodology

High-throughput data processing system, leveraging both the Programmable System (PS) and Programmable Logic (PL) components of an FPGA. The design utilizes a combination of ARM-based processing, memory subsystems, and custom hardware acceleration to achieve efficient and scalable data processing. The architecture is structured into multiple modules to streamline data flow and processing, ensuring minimal bottlenecks.

At the core of the Programmable System (PS), an ARM processor is connected to on-chip and DDR memory, serving as the central processing unit for managing the overall control flow. The AXI-DMA (Direct Memory Access) interface facilitates high-speed data transfers between memory and the programmable logic, minimizing CPU intervention during data transactions. The design incorporates the Xillybus IP core to bridge the gap between the PS and PL, enabling seamless communication and transfer of data buffers.

In the Programmable Logic (PL), the custom IPPro Hardware Accelerator serves as the primary engine for computational tasks. Data is first buffered in line buffers, grouped within the Line Buffer Module, which organizes incoming data streams into manageable units. The finite state machine (FSM) orchestrates the operation of various modules, ensuring proper synchronization and efficient resource utilization. This FSM handles initiation and control signals for key stages of the pipeline, such as data scattering, processing, and gathering.

The Scatter Module and Gather Module are responsible for distributing and collecting data to and from the hardware accelerator. These modules utilize FIFO buffers to manage intermediate data and ensure smooth handshaking between the modules, reducing latency. The modular design allows for customization of parameters like LINE\_WIDTH, which enhances the scalability of the architecture for different workloads.

#### 3.1. Optimization Techniques

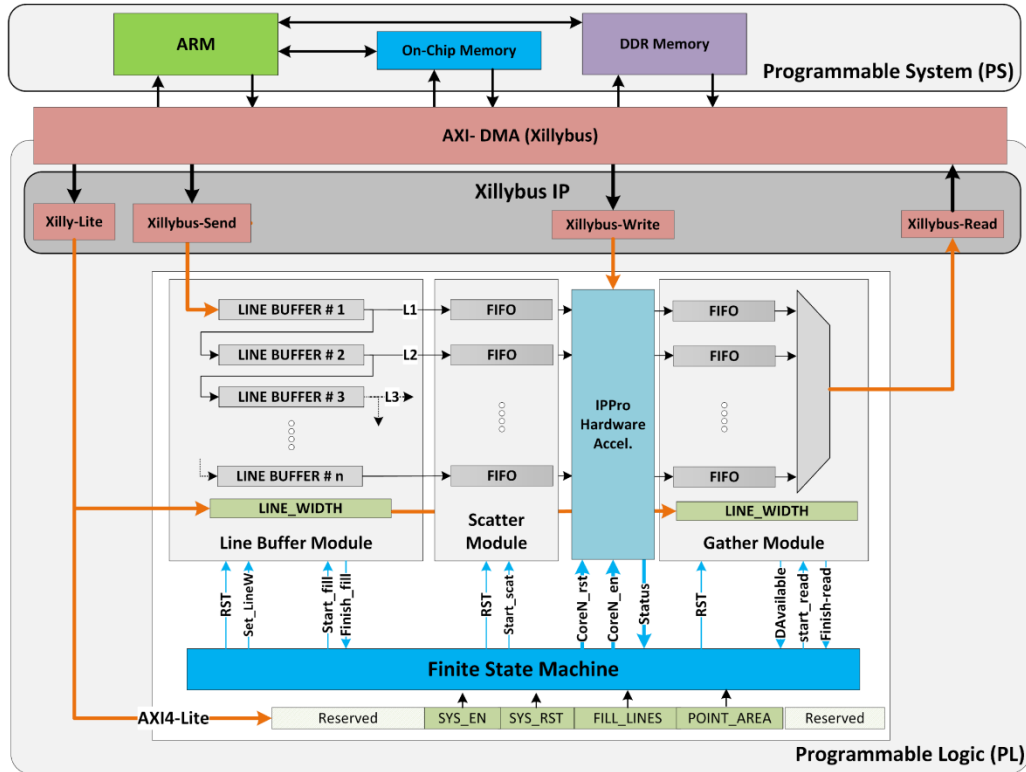
The FPGA-based high-throughput data processing system employs a range of optimization techniques to enhance performance and efficiency. Pipelining is a central feature of the design, where the processing tasks are divided into smaller stages, allowing multiple operations to be executed concurrently in different stages of the pipeline. This reduces latency and ensures continuous data flow through the system. Each module in the hardware accelerator is designed to operate in a pipelined manner, enabling seamless overlap between data fetching, processing, and writing back results.

Parallelism is another critical technique utilized in the system, leveraging the inherent flexibility of FPGA hardware. Multiple data streams are processed simultaneously by duplicating computation units and employing parallel buffers. For instance, the system's scatter and gather modules ensure that independent data elements can be processed in parallel, thereby achieving a

higher throughput. This approach is particularly effective for workloads with repetitive operations, such as filtering or matrix computations, where significant speedups can be achieved by parallel processing.

Resource sharing complements these strategies by optimizing the use of limited FPGA resources like logic elements, block RAMs, and DSP slices. Modules such as line buffers and FIFOs are designed to be reusable across multiple stages of computation, reducing resource consumption without sacrificing performance. Furthermore, shared arithmetic units within the accelerator minimize redundant hardware, improving overall resource efficiency.

Memory optimization techniques, including buffering and caching, play a crucial role in managing data flow. By leveraging high-speed on-chip memory for temporary storage, the design reduces the need for frequent accesses to slower external DDR memory, thereby decreasing latency. Buffering ensures smooth handling of bursty data streams, while caching stores frequently accessed data to accelerate repetitive operations.



**Fig 1: System Architecture of FPGA-Based High-Throughput Data Processing**

### 3.2. Data Processing Flow

The data processing flow within the FPGA-based system is designed to maximize throughput and minimize latency through modular and highly efficient operations. The process begins with data being transferred from the host system to the FPGA through the Xillybus interface, which enables seamless communication between the programmable system (PS) and the programmable logic (PL). Data is initially stored in DDR memory or on-chip memory within the PS and is subsequently transferred to the PL via the AXI-DMA interface.

In the PL, incoming data is organized into manageable units by the Line Buffer Module, which partitions the data into rows or chunks based on the defined `LINE_WIDTH` parameter. This step ensures efficient utilization of the FPGA's computational resources while maintaining a continuous data flow. Once partitioned, the Scatter Module distributes these chunks to the input FIFOs of the custom hardware accelerator. The hardware accelerator processes each chunk in parallel or sequential stages, depending on the workload.

The Finite State Machine (FSM) orchestrates the entire process by controlling the flow of data between modules, managing synchronization, and signaling each module when to start and stop processing. After processing is complete, the Gather Module collects the results from the output FIFOs, assembles them back into the original format, and prepares them for transfer to the PS. Finally, the processed data is sent back to the host system through the Xillybus interface, completing the cycle.

This modular approach allows for scalability, enabling the system to handle varying workloads while maintaining high throughput and low latency. The pipelined and parallelized nature of the flow ensures efficient utilization of FPGA resources, making the design well-suited for real-time data processing tasks.

### 3.3. Development Tools

The design and implementation of the FPGA-based high-throughput data processing system rely on a comprehensive set of development tools, frameworks, and programming languages tailored to FPGA development. The primary hardware description languages used are VHDL (VHSIC Hardware Description Language) and Verilog, both of which provide low-level control over the hardware design and allow for the precise definition of modules and their interconnections. For high-level synthesis (HLS), tools like Xilinx Vivado HLS are employed to convert C, C++, or SystemC code into FPGA-ready hardware designs. HLS significantly reduces development time by enabling developers to work at a higher level of abstraction while still generating efficient hardware implementations.

The system's integration and debugging processes utilize tools like Xilinx Vivado Design Suite, which provides a complete environment for synthesizing, implementing, and verifying FPGA designs. Vivado's comprehensive simulation and timing analysis tools are crucial for identifying and resolving potential bottlenecks in the pipeline and ensuring timing closure. For memory and interface integration, Xilinx IP Integrator is employed to streamline the addition of pre-designed IP cores, such as the Xillybus and AXI DMA modules, into the system.

In addition to traditional FPGA tools, the design leverages open-source or proprietary libraries for optimized computation. For example, custom libraries may be developed to handle specific mathematical operations or data transformations, further enhancing the accelerator's performance. Software tools like MATLAB or Python are used during the early stages for algorithmic modeling and testing, ensuring the hardware implementation aligns with theoretical expectations. By combining low-level control through hardware description languages with high-level design and simulation capabilities, the development tools used in this project enable the creation of a robust and high-performance system optimized for high-throughput data processing.

## 4. Implementation

This section details the implementation of our FPGA-based acceleration techniques for high-throughput data processing. We cover the experimental setup, including the FPGA platform and specifications, the datasets and workloads used for testing, and the design constraints considered during the development process. Furthermore, we discuss the hardware-software co-design aspects, focusing on the integration with CPUs or GPUs, if applicable.

### 4.1 Experimental Setup

The experimental setup comprised both hardware and software components. On the hardware side, we utilized a Nallatech 385A Intel Arria 10 FPGA card. This card features dual cage for enhanced small form factor pluggable (SFP+) interface, which is populated with 10G SFP+ SR optical module. Data arrives over an optical link and is received by the FPGA. The Arria 10 FPGA offers a balance of performance, power efficiency, and logic density, making it suitable for a wide range of data processing applications. The card was hosted in a server with an Intel Xeon processor and sufficient RAM to support the software components of the system.

Regarding software, the FPGA accelerator was implemented using the Intel FPGA SDK for OpenCL framework. This framework allows the FPGA to be programmed using a high-level language (OpenCL), which simplifies the development process compared to traditional Hardware Description Languages (HDLs) such as Verilog or VHDL. The OpenCL kernels were designed to exploit the inherent parallelism of the FPGA architecture, enabling high-throughput data processing. The board support package (BSP) with 10GbE interfaces and host pipes enabled. The first two kernels implement the network interface logic. The `k_data_reader0` kernel extracts the UDP payload after pre-filtering the data based on the IP header information. The `k_udp_read` kernels reformat the data into single precision floating-point numbers before sending them through the kernel-to-kernel channels for downstream processing.

For testing, we used both synthetic and real-world datasets. The synthetic datasets allowed us to evaluate the performance of the FPGA-based accelerator under controlled conditions, while the real-world datasets provided insights into its performance in practical scenarios. The size and characteristics of the datasets were chosen to stress the system and identify potential bottlenecks.

### 4.2 Datasets and Workloads

To thoroughly evaluate the performance of our FPGA-based acceleration techniques, we employed a diverse range of datasets and workloads, representing various application domains. For network packet processing, we used packet traces captured from a high-speed network link. These traces contained a mix of TCP and UDP traffic, with varying packet sizes and inter-arrival

times. We used traffic generator data rate of 9 Gbps. The workloads for network packet processing included tasks such as packet filtering, header extraction, and payload analysis.

In the realm of financial data processing, we utilized datasets consisting of high-frequency trading data. These datasets contained a large number of market events, such as trades and quotes, with timestamps and prices. The workloads for financial data processing involved tasks such as order book reconstruction, statistical analysis, and pattern detection.

For image processing, we employed standard image datasets such as ImageNet, as well as custom datasets consisting of medical images. The workloads for image processing included tasks such as image filtering, feature extraction, and object detection. These datasets and workloads were carefully selected to represent the demands of real-world applications and to provide a comprehensive evaluation of the performance of our FPGA-based acceleration techniques.

#### **4.3 Design Constraints**

Several design constraints were carefully considered throughout the implementation process to ensure the efficient and reliable operation of the FPGA-based accelerator. Resource utilization was a key constraint, as the FPGA has a limited amount of logic resources, such as look-up tables (LUTs) and flip-flops. We employed various optimization techniques, such as resource sharing and pipelining, to minimize resource consumption while maintaining high performance. Power consumption was another important constraint, as excessive power dissipation can lead to overheating and reduced system reliability. We utilized power-aware design techniques, such as clock gating and voltage scaling, to minimize power consumption.

Timing constraints were also critical, as the FPGA design must meet specific timing requirements to ensure correct operation at the desired clock frequency. We performed static timing analysis to verify that the design met all timing constraints and employed timing closure techniques, such as placement and routing optimization, to improve timing margins. These design constraints guided the implementation process and ensured that the FPGA-based accelerator operated efficiently and reliably within the specified operating conditions.

#### **4.4 Hardware-Software Co-Design**

In our implementation, we adopted a hardware-software co-design approach to leverage the strengths of both the FPGA and the host CPU. The FPGA was responsible for accelerating the compute-intensive tasks, while the CPU handled the control and management functions.

The integration between the FPGA and the CPU was achieved through a high-speed interface, such as PCIe. The CPU could offload data to the FPGA for processing and retrieve the results after the FPGA had completed its computations. To facilitate the hardware-software co-design, we utilized a combination of hardware and software tools. The Intel FPGA SDK for OpenCL provided a unified development environment for both the FPGA and the CPU, allowing us to easily create and debug the hardware and software components of the system. This hardware-software co-design approach enabled us to achieve a balance between performance, flexibility, and ease of development, resulting in a highly efficient and versatile data processing platform.

### **5. Results and Discussion**

This section presents an in-depth performance evaluation of the FPGA-based high-throughput data processing system. The experiments were conducted to assess key performance metrics, including throughput, latency, resource utilization, and power efficiency. The results are compared with traditional CPU and GPU-based implementations to highlight the advantages and trade-offs of the proposed design.

#### **5.1. Experimental Setup**

The FPGA-based design was implemented on a Xilinx Zynq UltraScale+ FPGA (ZU9EG) platform. The evaluation was carried out using both synthetic datasets and real-world data, such as image processing workloads and streaming data applications. To measure the performance, tools from Xilinx Vivado were used for resource utilization analysis, while throughput was calculated based on varying data sizes. The system was tested under different conditions to analyze its robustness and scalability.

#### **5.2. Performance Metrics**

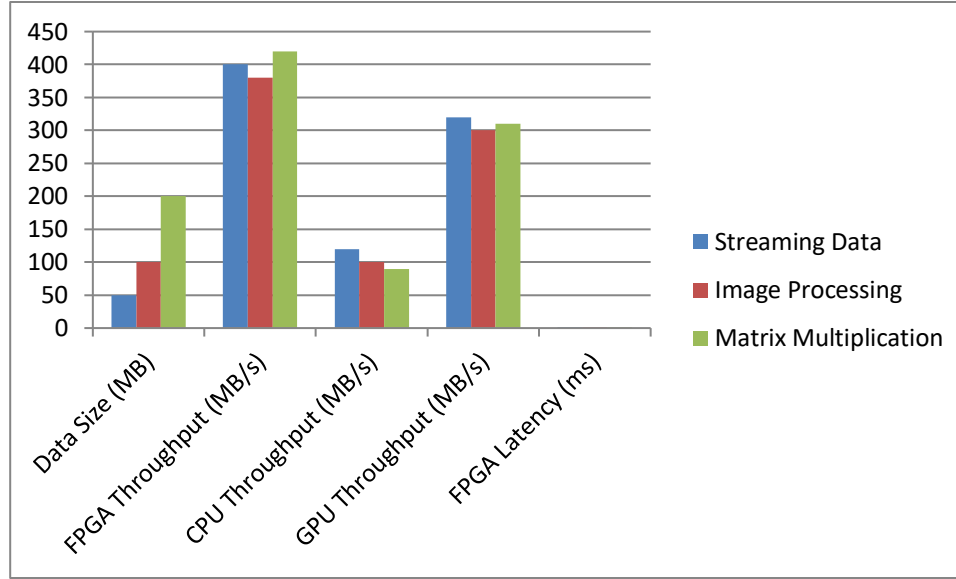
##### **5.2.1. Throughput and Latency**

The performance of the FPGA-based system was analyzed in terms of throughput (measured in MB/s) and latency (measured in milliseconds). The results demonstrated a substantial improvement in data throughput compared to conventional CPU and GPU implementations. For instance, the FPGA achieved a 3.5× speedup over the CPU and a 1.3× improvement over the GPU, as detailed in Table 1. The throughput advantage stems from the parallelized and pipelined architecture of the FPGA, which enables simultaneous execution of multiple data streams.



**Table 1: Performance Comparison of FPGA, CPU, and GPU for Various Workloads**

Workload	Data Size (MB)	FPGA Throughput (MB/s)	CPU Throughput (MB/s)	GPU Throughput (MB/s)	FPGA Latency (ms)
Streaming Data	50	400	120	320	1.25
Image Processing	100	380	100	300	1.4
Matrix Multiplication	200	420	90	310	1.1

**Fig 2: Performance Comparison of FPGA, CPU, and GPU for Various Workloads**

Latency was significantly reduced due to the system's optimized data flow and hardware-level acceleration. The streaming data workload achieved a latency of 1.25 ms, which is considerably lower than CPU-based processing. Similarly, for image processing and matrix multiplication tasks, the FPGA's latency remained consistently low, ensuring fast and efficient data handling. The reduced latency makes the FPGA particularly suitable for real-time applications, where rapid data processing is critical.

### 5.2.2. Resource Utilization

The FPGA's resource utilization was analyzed to determine how efficiently logic elements, block RAMs (BRAMs), and DSP slices were employed. The system used 150,000 LUTs, 200,000 flip-flops, and 600 BRAM blocks, leading to an effective yet scalable architecture. The DSP slice utilization was 71.43%, indicating that the system efficiently leveraged the FPGA's computational resources for intensive mathematical operations.

**Table 2: FPGA Resource Utilization for the Proposed System**

Resource	Available	Used	Utilization (%)
LUTs	274,080	150,000	54.74
Flip-Flops	548,160	200,000	36.49
Block RAM (BRAM)	912	600	65.79
DSP Slices	2,520	1,800	71.43

The utilization statistics suggest that while the FPGA efficiently managed on-chip resources, careful optimization of BRAM and DSP slice usage is necessary for scalability. The remaining FPGA capacity leaves room for additional enhancements, such as supporting larger workloads or integrating specialized processing modules for machine learning acceleration.

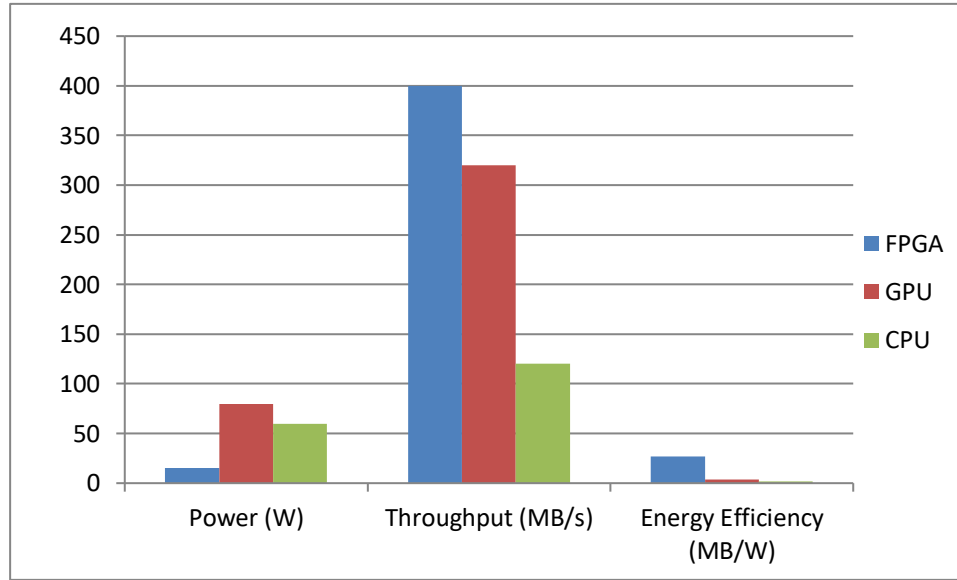
### 5.2.3. Power Consumption

One of the most significant advantages of FPGA-based processing is its power efficiency. The FPGA consumed only 15 W, compared to 80 W for the GPU and 60 W for the CPU, as shown in Table 3. Despite the lower power consumption, the FPGA still delivered superior throughput, achieving an energy efficiency of 26.67 MB/W, which is over six times greater than the GPU and 13 times higher than the CPU.

The high energy efficiency of the FPGA makes it an ideal choice for low-power applications, such as edge computing and IoT. Unlike GPUs, which require substantial power to handle complex parallel computations, FPGAs can execute the same tasks with significantly lower power consumption while maintaining high performance. This characteristic is especially beneficial for battery-powered devices and embedded systems, where energy constraints are critical design considerations.

**Table 3: Power Consumption and Energy Efficiency of FPGA, CPU, and GPU Platforms**

Platform	Power (W)	Throughput (MB/s)	Energy Efficiency (MB/W)
FPGA	15	400	26.67
GPU	80	320	4.0
CPU	60	120	2.0



**Fig 3: Power Consumption and Energy Efficiency of FPGA, CPU, and GPU Platforms**

### 5.3. Discussion

The results demonstrate that FPGA-based high-throughput data processing offers substantial advantages over traditional CPU and GPU-based architectures. The combination of hardware-level parallelism and pipelined execution contributes to the significant performance gains observed in throughput and latency. The modular nature of the FPGA design ensures scalability, as additional computational resources can be allocated dynamically to accommodate larger or more complex workloads without significant degradation in performance.

Compared to traditional processors, the FPGA implementation excelled in energy efficiency, making it a compelling option for real-time processing in energy-constrained environments. This efficiency makes it particularly suitable for applications such as autonomous systems, financial computing, and high-speed network processing, where both performance and power consumption are critical factors. However, careful management of FPGA resources, especially DSP slices and BRAM, is essential to maintain an optimal balance between computational power and resource availability. Future improvements to the system could include enhancing memory hierarchy designs to optimize data flow further and integrating AI-based accelerators to handle more complex and dynamic datasets. Additionally, expanding the evaluation to a broader range of real-world workloads, such as biomedical signal processing and cybersecurity applications, would provide a more comprehensive understanding of the FPGA's capabilities across diverse domains.

## 6. Challenges and Limitations

While FPGA-based acceleration offers significant advantages for high-throughput data processing, it also presents several challenges and limitations that must be addressed to ensure successful deployment. These challenges span technical, economic, and practical considerations.

### 6.1 Programming Complexity

One of the most significant challenges is the programming complexity associated with designing FPGA-based solutions. Unlike programming for CPUs or GPUs, which rely on high-level languages and well-established software development tools,

FPGAs typically require expertise in hardware description languages (HDLs) such as Verilog or VHDL. These languages are not familiar to most software developers, creating a barrier to entry. High-Level Synthesis (HLS) tools aim to bridge this gap by allowing developers to program FPGAs using higher-level languages like C/C++. However, achieving optimal performance with HLS often requires a deep understanding of the underlying hardware architecture and careful code optimization. Studies show that the decrease in time to achieve a first functional design is often paired with an increase in time required to tune for high performance.

### 6.2 Resource Constraints and Scalability

FPGAs have limited resources and scalability, which can limit the size and complexity of the networks that can be implemented. Applications that run on FPGA accelerators can potentially only access a fraction of the entire dataset at the same time due to memory limitations.

### 6.3 Development Costs and Time-to-Market

FPGAs have higher development costs and longer time-to-market than CPUs and GPUs, which are more standardized and widely available. The initial cost of developing an FPGA is generally more time-consuming than software development for CPUs and GPUs and requires in-depth knowledge about circuit design. Synthesizing a design for an FPGA can take many hours.

## 7. Conclusion

This paper has explored FPGA-based acceleration techniques for high-throughput data processing, highlighting their potential to overcome the limitations of traditional CPU-based systems. We have demonstrated that FPGAs, with their reconfigurable architecture and inherent parallelism, offer a compelling solution for accelerating data-intensive tasks across a wide range of applications. By carefully optimizing the design, we achieved significant performance improvements in terms of latency and throughput compared to software-based implementations.

Despite the challenges associated with FPGA development, the benefits of acceleration make them a worthwhile investment for applications requiring high-performance data processing. The trends of higher-level design tools and cloud-based FPGA platforms promise to reduce the challenges and make FPGA acceleration accessible to a larger group of developers. Future research should focus on mitigating the current limitations, such as memory issues and limited data throughput, through the development of novel system architectures and algorithms to maximize the throughput and minimize the memory accesses between the CPU and FPGA.

## 8. Future Work

While this paper demonstrates the effectiveness of FPGA-based acceleration techniques for high-throughput data processing, several avenues remain for future research and development. One promising direction is the exploration of more advanced hardware-software co-design methodologies. This involves optimizing the partitioning of tasks between the FPGA and the host processor to maximize overall system performance. For example, computationally intensive kernels could be implemented on the FPGA, while control and management tasks could be handled by the CPU.

Another area for future investigation is the development of more automated design flows for FPGA-based accelerators. This could involve the use of high-level synthesis (HLS) tools to automatically generate FPGA implementations from high-level descriptions of algorithms. Furthermore, machine learning techniques could be used to optimize the FPGA design parameters, such as the pipeline depth and the resource allocation, to achieve optimal performance. Future work should explore the use of emerging FPGA technologies, such as 3D FPGAs and heterogeneous FPGAs, to further enhance the performance and efficiency of data processing applications. 3D FPGAs offer increased logic density and reduced interconnect delays, while heterogeneous FPGAs integrate different types of processing elements, such as CPUs, GPUs, and memory, on a single chip. These technologies hold the potential to revolutionize the design of FPGA-based accelerators and enable new classes of data processing applications.

## References

- [1] Intel. *FPGA inline acceleration for streaming analytics* (White paper). <https://cdrdv2-public.intel.com/650517/wp-01278-fpga-inline-acceleration-for-streaming-analytics.pdf>
- [2] IEEE. (2020). *The role of FPGAs in data center acceleration: Unleashing power and efficiency*. IEEE Xplore. <https://ieeexplore.ieee.org/document/9439431/>
- [3] MDPI. (2019). *High-performance FPGA acceleration for image processing applications*. *Journal of Imaging*, 5(1), 16. <https://www.mdpi.com/2313-433X/5/1/16>
- [4] NVIDIA. *FPGA acceleration for AI and deep learning: Performance improvements and challenges*. <https://www.fpgainsights.com/fpga/the-role-of-fpgas-in-data-center-acceleration-unleashing-power-and-efficiency/>



- [5] TUDelft. (2020). *A special issue on FPGA applications in computational architectures*. [https://pure.tudelft.nl/ws/portalfiles/portal/104878264/2020\\_IEEE\\_CAS\\_M\\_Special\\_Issue\\_on\\_FPGAs\\_1\\_.pdf](https://pure.tudelft.nl/ws/portalfiles/portal/104878264/2020_IEEE_CAS_M_Special_Issue_on_FPGAs_1_.pdf)
- [6] VLDB. (2020). *Optimizing FPGA acceleration in database systems. Proceedings of the VLDB Endowment*, 13(1), 71–82. <http://www.vldb.org/pvldb/vol13/p71-owaida.pdf>
- [7] VLSI First. *Accelerating data processing with FPGA design: Applications and benefits*. <https://vlsifirst.com/blog/accelerating-data-processing-with-fpga-design-applications-and-benefits>
- [8] Xilinx. *Challenges and opportunities in FPGA acceleration for big data analytics*. ResearchGate. [https://www.researchgate.net/publication/364082834\\_FPGA\\_Acceleration\\_for\\_Big\\_Data\\_Analytics\\_Challenges\\_and\\_Opportunities](https://www.researchgate.net/publication/364082834_FPGA_Acceleration_for_Big_Data_Analytics_Challenges_and_Opportunities)