



Digital Circuit Acceleration Techniques for Computer Vision Tasks

Boyang Wu

School of Electronic Information, Northeast University, Qinghuangdao, China

202212644@stu.neuq.edu.cn

Abstract. Digital circuit acceleration is pivotal in enhancing the performance and efficiency of modern electronic systems, particularly in applications such as real-time processing, low-power embedded devices, and high-speed data transmission. This article explores diverse methodologies to accelerate digital circuits, focusing on four core approaches. First, a Field-Programmable Gate Array (FPGA)-based functional verification framework is introduced. It leverages hardware-software co-verification and dynamic coverage analysis to achieve an 1132 \times acceleration in verification cycles. Second, a reconfigurable architecture for regular expression engines is proposed, combining deterministic state machines and instruction-driven designs to optimize throughput and resource utilization. Third, a low-power multimode digital front-end circuit is developed using dynamic CIC filter configurations and Farrow structure optimizations, reducing power consumption by 11.9%, particularly in GSM and TD-SCDMA applications. Lastly, Tiny YOLO hardware acceleration on Zynq SoCs employs network pruning and pipeline parallelism, enabling real-time vehicle detection at 24 FPS, which is vital for autonomous driving systems. Additionally, computer vision architectures and SVM-based predictive models are analyzed for their roles in system optimization. These advancements collectively address critical challenges in speed, energy efficiency, and adaptability, offering significant value for both academic research and industrial applications in embedded systems and beyond.

Keywords: FPGA, Hardware Acceleration, Low-Power Circuits

1 Introduction

Digital circuit acceleration has become indispensable in meeting the escalating demands for high-speed computation, energy efficiency, and real-time responsiveness across industries such as telecommunications, automotive systems, and industrial automation [1-3]. The evolution of semiconductor technology and the proliferation of complex algorithms, such as deep learning and signal processing, have intensified the need for optimized circuit designs. For instance, traditional software-based verification methods for Application-Specific Integrated Circuits (ASICs) often suffer from prolonged cycle times, necessitating hardware-assisted solutions like FPGA-

© The Author(s) 2025

A. J. Moshayedi (ed.), *Proceedings of the 2025 2nd International Conference on Electrical Engineering and Intelligent Control (EEIC 2025)*, Advances in Engineering Research 279,

https://doi.org/10.2991/978-94-6463-864-6_16

based co-verification [4]. Similarly, the rise of 5G and IoT has driven innovations in low-power circuit architectures, such as dynamic reconfiguration techniques, to support multi-standard compatibility [5].

Recent advancements in reliability assessment methodologies further underscore the importance of robust circuit design. For example, Chen et al. proposed a hybrid approach combining Spice reliability simulations and High-Temperature Operating Life (HTOL) experiments to evaluate aging effects in SoC digital circuits [6]. Their work established a Power-Law-based lifetime prediction model, validated under stress voltage and time conditions, which revealed that device degradation aligns with the R-D model. These insights are critical for ensuring long-term reliability in accelerated systems, particularly under harsh operational environments.

Moreover, practical applications of digital acceleration extend to embedded systems with stringent real-time requirements. Ye et al. designed an intelligent medicine delivery vehicle using STM32 and OpenMV-based vision systems, achieving automated line tracking and digital ward recognition through normalized cross-correlation and contour extraction algorithms [7]. This application highlights the necessity of efficient hardware-software co-design, where optimized digital circuits enable rapid image processing (e.g., 30 fps acquisition) and low-latency control (e.g., <10-second system delay), crucial for mission-critical healthcare scenarios.

This article systematically examines four key acceleration strategies. First, FPGA-based verification frameworks integrate Hard Processor Systems (HPS) and protocol bridges to enable real-time signal analysis, reducing verification cycles from over 200,000 seconds to 179 seconds [8]. Second, reconfigurable regular expression engines employ state machine optimizations and instruction-driven architectures to achieve 6.4 Gbps throughput with 90% resource savings. Third, multimode digital front-end circuits utilize dynamic filter configurations and structural reconfiguration to enhance power efficiency by 11.9% while supporting GSM and TD-SCDMA standards. Finally, Tiny YOLO acceleration on Zynq SoCs demonstrates how network pruning and hardware-software co-design enable real-time object detection at 24 FPS, critical for autonomous driving applications. Complementary analyses of computer vision architectures and predictive models further underscore the interdisciplinary nature of circuit optimization. These methodologies collectively address the dual challenges of performance and sustainability, paving the way for next-generation embedded systems.

2 Digital Circuit Acceleration Techniques

2.1 FPGA-Based Digital Circuit Functional Verification

The utilization of HPS (Hard Processor System) is proposed to accelerate ASIC functional verification, with core contributions including a hardware-software co-verification framework that connects HPS and FPGA logic via an AXI-to-DUT protocol bridge. In this framework, the Cortex-A CPU generates C/C++ verification stimuli, while the ELA embedded logic analyzer captures signals in real-time,

reducing verification cycle time from 203,572 seconds in software simulation to 179 seconds, achieving an 1132x acceleration. Cross-clock-domain data transfer is enabled through asynchronous FIFOs and synchronizers in the protocol bridge, supporting DUT clock (98.6MHz) and AXI clock domain isolation, with a bandwidth of 25.2Gb/s—25x faster than traditional Ethernet. Additionally, dynamic analysis of functional coverage is achieved by having the HPS compare DUT outputs with a reference model in real-time, generate coverage reports, support breakpoint debugging and single-step tracking, and consume stable logic resources when multiplexing verification IP. In a real test on the Cyclone V SoC platform, the validation of a 68.2kLE image processing system demonstrated protocol bridge latency of less than 10ns and frequency convergence better than 54.8% of traditional FPGA prototype validation at 63.5MHz [1].

2.2 Reconfigurable Architecture Design for Regular Expression Hardware Acceleration Engine

This method proposes two high-speed regular matching scheme. State Machine Direct Connect Architecture: converts NFA to deterministic state machine, supports hard-wired implementation of base metacharacters (e.g., [a-z], a {n,m}). State transfer logic is described through Verilog, reaching 800MHz clock frequency on Cyclone IV FPGA with throughput rate of 6.4Gbps (8bit/cycle), which is 1-2 times faster than similar solutions, but requires customized circuits for each expression.

Instruction-driven reconfigurable architecture: 78-bit instruction format is designed to store parameters such as matching mode (single character/range) and count type (fixed/range/greedy) through RAM, and reuse 3 comparators ($=$, \geq , \leq) to realize "double reuse". Adopting AXI bus interaction and data caching strategy, the delay in matching trigger is 5-8 cycles, and the delay in untriggered trigger is 1 cycle, and utilizing only 355 LE, representing a 90% reduction compared to traditional methods. Experiments show that the reconfigurable scheme supports complex expressions with a matching success rate of 100%, and the resource consumption is 37% of BRAM 4.9Mb [2].

2.3 Low-Power Optimized Design of Multimode Digital Front-End Acceleration Circuitry

Aiming at the demand for software radio multi-standard compatibility, this paper proposes a digital front-end optimization scheme based on a CIC filter and Farrow structure. Core innovations include the CIC Filter Dynamic Order Configuration, which adjusts the CIC order based on the oversampling rate (OSR), reducing the required orders for 50dB stopband attenuation from 8 to 2-4 for $OSR \geq 4$, thereby decreasing computation by 29.2%. The Farrow filter is optimized to support low OSR scenarios by adjusting 3rd order polynomial coefficients through least squares fitting, extending its application to 2x OSR scenarios with 0.1dB passband ripple, 50dB stopband attenuation, a 50% reduction in operating frequency, and the elimination of subsequent half-band filters. Additionally, a multi-level structural reconfiguration is

employed, adopting an NRCIC-Farrow-FIR cascade architecture (Eq. (12)) to support multi-standards such as GSM and TD-SCDMA, achieving an EVM index of 1.51%/1.19% and reducing power consumption by 11.9% compared to traditional schemes. The proposed design also achieves an 8:1 storage compression ratio and supports 2-32 times configurable rate conversion, with the object to be inspected automatically conveyed to the inspection area by the feeding device shown in Figure 1 [3].




Fig. 1. Structure of shipping device network [3]

2.4 Tiny YOLO Vehicle Detection on Zynq SoC with Hardware Acceleration

An X-TINY YOLO algorithm optimized for the Zynq-7020 SoC is proposed to achieve real-time vehicle detection by co-designing the compressed network structure with hardware resources. Core contributions include sparse network pruning, which dynamically adjusts convolutional kernels (e.g., reducing the second-layer kernel from 32 to 24). Combined with batch-normalized parameter fusion, this reduces computation volume by 48.6% while maintaining a 62% mAP. Data chunking transmission and computation optimization are implemented by designing a weight matrix and feature map chunking strategy to address the DSP48E (220) resource limitation at the PL side. Using the AXI_ACP interface and DMA_SG transmission, single-layer convolutional computations are decomposed into multiple steps (e.g., the first layer is divided into 56 steps), achieving BRAM utilization of 52% and DSP utilization of 50%. A multi-stage pipeline parallel architecture is deployed with an efficient pipeline ($\Pi=1$) at the PL side, combined with a dynamic data zero-completion strategy, achieving 24 FPS at 11.3W power consumption, which is 4.7 times faster than a GPU solution. Experiments demonstrate that the XC7Z020 chip resource consumption is 50% for DSP and 52% for BRAM, achieving real-time performance at 640×480 resolution and reducing hardware requirements by 50% compared to the original TINY YOLO hardware design [4].

2.5 Core Applications and Limitations of Deep Learning Architectures in Computer Vision Analyzed

This study systematically analyzes the application of Convolutional Neural Networks (CNN), Restricted Boltzmann Machines (RBM) and Deep Belief Networks (DBN) in computer vision. The authors elaborate the local sense field, weight sharing and spatial downsampling mechanisms of CNN through mathematical modeling, point out its advantage of reducing the number of parameters through sparse weight matrix, and implement multi-scale feature fusion combining FPN top-down semantic information transfer and PANet bottom-up enhancement of localization features in YOLOv5. RBM improves training efficiency through the two-part graph structure and unsupervised contrast scattering algorithm, while DBN faces high computational cost (e.g., ignoring the 2D structure of images) due to the layer-by-layer greedy training strategy. Efficiency, while DBN faces the drawback of high computational cost (e.g., ignoring the 2D structure of images) due to the layer-by-layer greedy training strategy. The study emphasizes the utility of CNNs in target detection (e.g., Faster R-CNN) and face recognition (e.g., FaceNet's 3D alignment technique), but lacks specific quantified accuracy metrics to support the claimed improvements and only points out the robustness improvement of YOLOv5 in feature fusion [5].

2.6 Design and Validation of Multimodal Perception-Driven Postoperative Intelligent Monitoring System for Ophthalmology

The system is deployed based on .NET 8 framework with Docker containers and integrates a smart camera (30 fps acquisition) and a fiber optic sensing pillow (non-invasive blood pressure monitoring). Emotion analysis was performed using a CNN model trained on the Fer2013 dataset (input 48×48 grayscale map) containing 3 layers of convolution (64-128-256 channels) and 3 layers of full connectivity (Dropout=0.5), outputting 7 classes of expression probabilities. Experiments show that the accuracy of emotion recognition in a single-person ward reaches 70%, though occlusion increases the false alarm rate to 25% in multi-person environments. In hardware, Zynq-7020 SoC realizes efficient data transfer from PS-PL side through AXI_ACP interface, and the overall delay of the system is controlled within 10 seconds, and the error of physiological parameter monitoring is $<2\%$ [8].

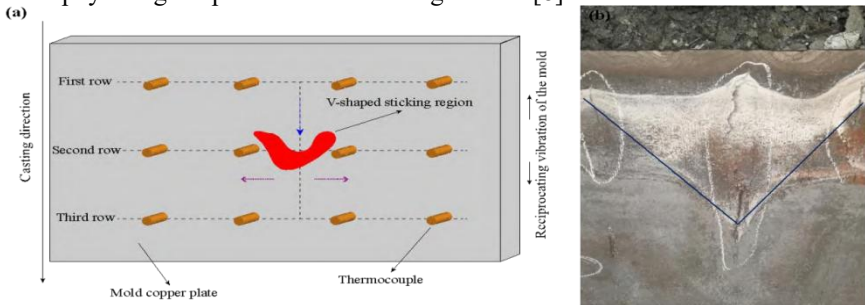


Fig. 2. Schematic diagrams of the "V-shaped" Sticking region and its two-dimensional plane expansion [8]

A schematic diagram of the V-shaped bonding region and its 2D planar extension is shown in Figure 2. Figure 2(a) shows the distribution characteristics of the V-shaped bonding region formed on the surface of the copper plate along the transverse and longitudinal expansion. The bonding region will continue to move longitudinally along the casting direction with the casting process, and at the same time, when the movement of the crystallizer itself and the billet shell movement in the opposite direction will be formed in the billet shell at the weak point of the tear, the tear in the repeated vibration of the crystallizer to further extend to the sides, and ultimately in the formation of the copper plate is similar to the "V" shaped extension of the area (labeled by the red area). Figure 2(b) visualizes the extension of the bonding area observed in Figure 2(a).

2.7 High-Precision Prediction Model for Continuous Casting Steel Leakage based on Feature Engineering and Svm

The thermocouple temperature signal of the crystallizer copper plate is mapped into HSV color space to create a two-dimensional thermogram. Additionally, a temperature change rate thermogram is extracted using a time-difference method with a 5-second interval. The Fourier descriptor is used to quantify the contour shape of the bonding region, and ten-dimensional feature vectors are constructed by combining the dynamic expansion factor and the center of gravity velocity ($v_x, v_y, v_{v_x}, v_{v_y}$). The RBF kernel function is selected for the SVM model ($C=100, C=100$, category weight 10:1), and seven key are screened out by sequence backward selection features (e.g., $Rave, S_E, FDRave, S_E, FD$), achieving 100% recall, a 9.93% false alarm rate, and an AUC of 0.98. (10 cases of missing steel + 131 normal). In contrast, random forest (100 trees) misses 2 cases due to feature redundancy (5-dimensional features).cases, which verifies the advantage of SVM under high dimensional small sample [9].

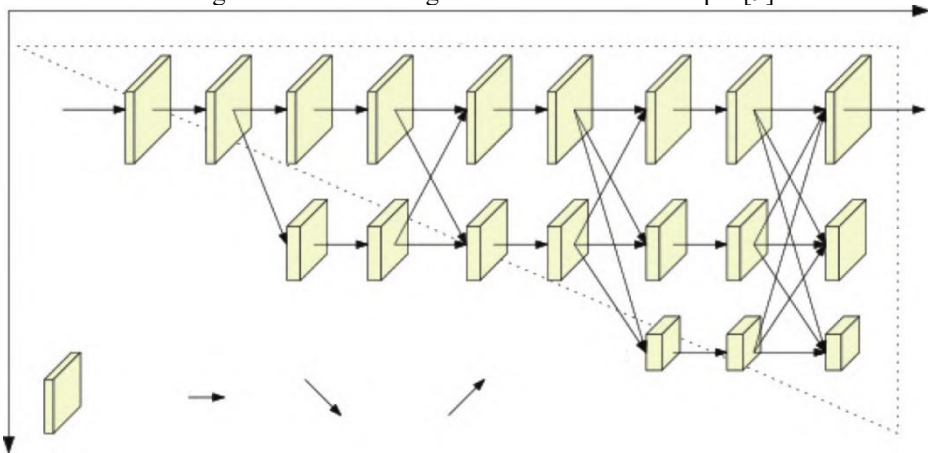


Fig. 3. Full-body key points detection framework network [9]

The full-body keypoint detection model shown in Figure 3 based on the HRNet backbone network fully utilises the advantages of multi-resolution feature fusion,

effectively supporting daily human posture detection tasks and achieving leading detection accuracy. Through end-to-end rapid inference, the model quickly and accurately outputs human body keypoint information from images. It leverages multi-resolution feature fusion to effectively support daily human posture detection tasks.

2.8 Tiny YOLO Hardware Acceleration and Resource Optimization for Embedded Scenarios

X-TINY YOLO reduces computation by 50% through kernel clipping (e.g., reducing the second layer from 32 to 24) and BN layer fusion. It implements parallel multiply-add operations on Zynq-7020 using a DSP48E array. The data chunking transfer strategy splits the weight matrix by 144 columns, combining the AXI_ACP interface and the DMA_SG mode to reduce the BRAM occupancy (148/280 chunks are required for a single layer). Tests show that the model achieves 62% mAP on the KITTI dataset, achieving 24 FPS with a 33.68 ms/frame elapsed time. And a resource footprint of only 50% DSP and 52% BRAM. Although the accuracy is slightly lower than that of the original TINY YOLO (67% mAP), the energy-efficiency ratio (24 FPS @ 142 MHz) offers a cost-effective option for embedded applications [10].

3 Conclusion

This article comprehensively reviews cutting-edge techniques for accelerating digital circuits, emphasizing their transformative potential in addressing speed, power, and adaptability constraints. The FPGA-based co-verification framework drastically reduces verification time through hardware-software synergy, while reconfigurable architectures for regular expression engines achieve unprecedented throughput-resource trade-offs. Low-power front-end designs exemplify how dynamic filter adjustments and structural reconfiguration minimize energy consumption, particularly in GSM and TD-SCDMA systems, without compromising functionality. Furthermore, Tiny YOLO optimizations on embedded platforms highlight the viability of lightweight deep learning models for real-time applications. Looking ahead, the integration of AI-driven automation into circuit design, the adoption of advanced process nodes (e.g., 3nm FinFET), and the exploration of quantum-inspired architectures are poised to redefine the boundaries of digital acceleration. These advancements will not only enhance existing systems but also unlock new possibilities in edge computing, smart manufacturing, and beyond, solidifying the role of digital circuits as the backbone of technological innovation. These advancements will enhance existing systems and unlock new possibilities in edge computing, such as real-time IoT analytics, and smart manufacturing, like predictive maintenance, solidifying digital circuits as the backbone of innovation.

References

1. Liu, X.Q., Yuan, G.S., and Qiao, S.S., Electron, J.: Accelerating Functional Verification for Digital Circuit with FPGA Hard Processor System, *Journal of Electronics & Information Technology*, 41(5), 1251–1256 (2019).
2. Tan, Z.Z.: Hunan University, Reconfigurable Design of Regular Expression Hardware Acceleration Engine (2021).
3. Zhao, Y. and Zhou, X.F.: Optimized wireless multi-mode digital front-end acceleration circuit design method, *Journal of China Computer System*, 34(8), 1940–1944 (2013).
4. Zhang, Y.K. and Liu, D.: Real-time implementation of improved TINY YOLO vehicle detection algorithm based on Zynq SoC hardware acceleration, *Journal of Computer Applications*, 39(1), 192–198 (2019).
5. Bai, J.: *Heilongjiang Science* 16, 0149–0151 (2025).
6. Ye, T. et al.: Automatic Medicine Delivery Vehicle Based on Vision System, *Shanxi Electronic Technology*, 23(1), 66–69 (2023).
7. Zhang, M. and Hong, Y.: Chin. Intelligent Monitoring System Based on Computer Vision and Artificial Intelligence, *Chinese Journal of Medical Instrumentation*, 49, 0074–0079 (2025).
8. Y. Wang et al.: Mold Breakout Prediction Model Based on Computer Vision and Machine Learning *Iron Steel* 59, 42–53 (2025)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

